

ON THE WITHIN-FAMILY KULLBACK-LEIBLER RISK IN GAUSSIAN PREDICTIVE MODELS[†]

BY GOURAB MUKHERJEE AND IAIN M. JOHNSTONE

Stanford University

We consider estimating the predictive density under Kullback-Leibler loss in a high-dimensional Gaussian model. Decision theoretic properties of the within-family prediction error – the minimal risk among estimates in the class \mathcal{G} of all Gaussian densities are discussed. We show that in sparse models, the class \mathcal{G} is minimax sub-optimal. We produce asymptotically sharp upper and lower bounds on the within-family prediction errors for various subfamilies of \mathcal{G} . Under mild regularity conditions, in the sub-family where the covariance structure is represented by a single data dependent parameter $\hat{\Sigma} = \hat{d} \cdot I$, the Kullback-Leibler risk has a tractable decomposition which can be subsequently minimized to yield optimally flattened predictive density estimates. The optimal predictive risk can be explicitly expressed in terms of the corresponding mean square error of the location estimate, and so, the role of shrinkage in the predictive regime can be determined based on point estimation theory results. Our results demonstrate that some of the decision theoretic parallels between predictive density estimation and point estimation regimes can be explained by second moment based concentration properties of the quadratic loss.

1. Introduction and main result.

1.1. Background. We consider a prediction set-up where the observed past data \mathbf{X} and the unobserved future data \mathbf{Y} are generated from a joint parametric density $f_{\boldsymbol{\theta}}(\mathbf{x}, \mathbf{y})$ where $\boldsymbol{\theta}$ is the unknown parameter. A perspective in prediction analysis is to use the concept of predictive likelihoods (Hinkley, 1979, Lauritzen, 1974) and its variants (Bjørnstad, 1990), to infer about the future \mathbf{Y} based on \mathbf{X} , with $\boldsymbol{\theta}$ playing the role of a nuisance parameter. Most predictive likelihoods (Butler, 1986) are functions of the future conditional density $f_{\boldsymbol{\theta}}(\mathbf{y} | \mathbf{X} = \mathbf{x})$ which is also referred to as the predictive

[†]This paper is based on Chapter 3 of G.M.'s Ph.D. thesis (to be submitted) and was titled as 'On efficient quadratic approximations of the predictive log-likelihood in a Gaussian sequence model' in previous preprints.

AMS 2000 subject classifications: Primary 62C20; Secondary 62M20, 60G25, 91G70

Keywords and phrases: prediction, predictive density, risk diversification, shrinkage, sparsity, high-dimensional, Kullback-Leibler, moment-based concentration inequalities

density (Geisser, 1971). Efficient estimates of the predictive density will ensure good predictive performances. Here, we study the problem of predictive density estimation in a high-dimensional Gaussian model.

We consider the multiple regression model analyzed in George and Xu (2008). Suppose, the observed past \mathbf{X} is independently generated from m_1 -dimensional product Gaussian density $N(A\boldsymbol{\theta}, \sigma_p^2 I)$ indexed by an n -dimensional unknown parameter $\boldsymbol{\theta}$ and known variance σ_p^2 and known $m_1 \times n$ data-matrix A . The future Y is generated from the m_2 -dimensional Gaussian density $N(B\boldsymbol{\theta}, \sigma_f^2 I)$ with the time-invariant parameter $\boldsymbol{\theta}$ and known $m_2 \times n$ dimensional matrix B and known future volatility σ_f^2 .

Homoscedastic Gaussian Predictive Model

$$\text{M.1} \quad \mathbf{X} \sim N(A\boldsymbol{\theta}, \sigma_p^2 I) \quad \text{and} \quad \mathbf{Y} \sim N(B\boldsymbol{\theta}, \sigma_f^2 I)$$

The location structure depend on the time-invariant unknown vector $\boldsymbol{\theta}$ of length n . If $\boldsymbol{\theta}$ is fixed the true predictive density of \mathbf{Y} would be $p_{(\boldsymbol{\theta}, B)}(\cdot) = N(B\boldsymbol{\theta}, \sigma_f^2 I)$. We would like to estimate it by density estimates $\hat{p}(\cdot | \mathbf{X} = \mathbf{x})$.

We use the information theoretic measure of Kullback and Leibler (1951) as the goodness of fit measure between the true and estimated distributions

$$\mathbf{L}(\boldsymbol{\theta}, \hat{p}(\cdot | \mathbf{x})) = \int p_{(\boldsymbol{\theta}, B)}(\mathbf{y}) \log \left(\frac{p_{(\boldsymbol{\theta}, B)}(\mathbf{y})}{\hat{p}(\mathbf{y} | \mathbf{x})} \right) d\mathbf{y}.$$

Averaging over the past observations \mathbf{X} , the predictive risk of the density estimate $\hat{p}(\cdot | \mathbf{X} = \mathbf{x})$ at $\boldsymbol{\theta}$ is given by

$$(1) \quad \boldsymbol{\rho}(\boldsymbol{\theta}, \hat{p}) = \iint p_{(\boldsymbol{\theta}, A)}(\mathbf{x}) p_{(\boldsymbol{\theta}, B)}(\mathbf{y}) \log \left(\frac{p_{(\boldsymbol{\theta}, B)}(\mathbf{y})}{\hat{p}(\mathbf{y} | \mathbf{x})} \right) d\mathbf{y} d\mathbf{x}.$$

The relative entropy predictive risk $\boldsymbol{\rho}(\boldsymbol{\theta}, \hat{p})$ measures the exponential rate of divergence of the joint likelihood ratio over a large number of independent trials (Larimore, 1983). The minimal predictive risk estimate maximizes the expected growth rate in repeated investment scenarios (Cover and Thomas, 1991, Chapter 6 and 15). Competitive optimal predictive schemes (Bell and Cover, 1980) for gambling, sports betting, portfolio selection, etc can be constructed from predictive density estimates with optimal Kullback-Leibler (KL) risk properties. Our Gaussian predictive framework can accommodate a fairly large number of prediction scenarios as often, in high-dimensional models, good normalization transformations of the data are available (Efron, 2010,

Chapter 1, Page 8). In the data compression set-up $\mathbf{L}(\boldsymbol{\theta}, \hat{p}(\cdot | \mathbf{x}))$ reflects the excess average code length that we need in Gaussian channels if we use the conditional density estimate \hat{p} instead of the true density to construct a uniquely decodable code for the data \mathbf{Y} given the past \mathbf{x} (McMillan, 1956). The notion can be extended to a sequential framework where minimizing the predictive risk would result in the minimum description length (Barron, Rissanen and Yu, 1998, Rissanen, 1984) based estimate of the true parametric density (Liang and Barron, 2005).

Here, we discuss efficient estimators in the class \mathcal{G}_n of all n -dimensional Gaussian distributions with positive definite (p.d.) covariances as the dimension increases, i.e.,

$$\mathcal{G} = \left\{ g : \mathbb{R}^n \rightarrow \mathbb{R}^+ \text{ such that } g = N(\mu, \Sigma) \text{ where } \mu \in \mathbb{R}^n \text{ and } \Sigma \text{ p.d.} \right\}.$$

For any class \mathcal{C} we define the predictive risk of the class as

$$\rho_{\mathcal{C}}(\boldsymbol{\theta}) = \inf_{\hat{p} \in \mathcal{C}} \rho(\boldsymbol{\theta}, \hat{p}).$$

As the true parametric density is also Gaussian, $\rho_{\mathcal{G}}(\boldsymbol{\theta})$ represents the within-family predictive risk. We also evaluate the predictive risk of the sub-family $\mathcal{G}[p]$ which contains all product Gaussian densities. We also make inferences in sparsity restricted parameter spaces. We impose an ℓ_0 constraint on the parameter space:

$$(2) \quad \Theta(n, s) = \left\{ \boldsymbol{\theta} \in \mathbb{R}^n : \sum_{i=1}^n \mathbb{I}[\theta_i \neq 0] \leq s \right\}.$$

This notion of sparsity is widely used in modeling highly interactive systems (represented by a large number of related parameters) which are dominated by only few significant effects. Sparse models have been successfully employed in biological sciences (Tibshirani et al., 2002), engineering applications (Donoho, 2006) and financial modeling (Brodie et al., 2009). The predictive model \mathbf{M} with ℓ_0 constraint on the location structure can be used for sparse coding and for prediction in sparse networks.

As in point estimation, risk calculations in \mathbf{M} would intrinsically depend on risk calculations in the orthogonal model:

Orthogonal Gaussian Predictive Model

$$\mathbf{M.2} \quad \mathbf{X} \sim N(\boldsymbol{\theta}, \sigma_p^2 I) \quad \text{and} \quad \mathbf{Y} \sim N(\boldsymbol{\theta}, \sigma_f^2 I)$$

where \mathbf{X} and \mathbf{Y} are both n – dimensional vectors. Most of our calculations will be in high-dimensions (which means $n \rightarrow \infty$ in the orthogonal model) though dimension independent bound will also be provided. As $n \rightarrow \infty$, **M.2** represents the Gaussian sequence model (Nussbaum, 1996) and has been widely studied in the function estimation framework (Johnstone, 2012). Estimation in **M.1** can be linked with the decision theoretic results in **M.2** through the procedure outlined in Donoho, Johnstone and Montanari (2011).

Our Contributions. Efficacy of predictive density estimates has been a subject of considerable interest in predictive inference. Aitchison (1975), Aslan (2006), Hartigan (1998), Komaki (1996) determined asymptotically optimal (admissible) Bayes predictive density estimates in fixed dimensional parametric family whereas minimax optimality in restricted parameter spaces has been discussed in Fourdrinier et al. (2011) and Kubokawa et al. (2012). Recently, Brown, George and Xu (2008), George, Liang and Xu (2006), Ghosh, Mergel and Datta (2008) extended the admissibility results to high dimensional Gaussian models. However, the optimal estimates are not necessarily Gaussian and using them in high-dimensional problems would involve computationally intensive methods. Here, we find optimal predictive density estimates within the Gaussian family and also compute their predictive risk. It is computationally easier to construct predictive attributes based on our optimal Gaussian predictive density estimates and the optimal Gaussian predictive risk assures guaranteed performances of our strategies.

Minimizing the Gaussian predictive risk involves simultaneous estimation of the location and scale parameters. The issue of joint estimation of location and scale (and to a degree the shape) has not been addressed before in one sample Gaussian models. However, separate estimation of location (Tibshirani, 2011) and covariance (Friedman, Hastie and Tibshirani, 2008) are well-studied topics in constrained Gaussian estimation. Also, as reviewed in George, Liang and Xu (2012) decision theoretic parallels exist between point estimation theory under quadratic loss and predictive density estimation under Kullback-leibler loss in high-dimensional Gaussian models. Here, our results demonstrate that some of these decision theoretic parallels (in the class \mathcal{G}) can be explained by second moment based concentration properties on the quadratic loss of location point estimators in high dimensions. The moment based approach used here for estimating the scale parameter bears resemblance to concepts seen elsewhere in prediction theory, particularly in the theory of cross validation (Yang, 2007) and covariance penalties for model selection (Efron, 2004, Ye, 1998) .

1.2. *Description of the main results.* We describe results of two different flavors concerning (i) the minimax risk of the class \mathcal{G} of Gaussian density estimates in sparsity restricted spaces (ii) the optimality (asymptotic admissibility, oracle inequality, risk upper bounds) of estimates in class \mathcal{G} in unrestricted space (over \mathbb{R}^n as $n \rightarrow \infty$) where the role of shrinkage comes into play. In order to describe the results, we need to introduce the following notations.

Notation and Preliminaries. As some of our results are dimension dependent, henceforth we refrain from using bold representation for vectors and denote the dimension in the subscript. Given any fixed sequence θ_∞ we represent the first n values by the n -dimensional vector θ_n whereas $\theta(n)$ denotes the n^{th} value, i.e $\theta_{n+1} = (\theta_n, \theta(n+1))$. By $\mathcal{G}_n[p]$ we denote the class of all n -dimensional product Gaussian densities

$$\mathcal{G}_n[p] = \left\{ g[\mu_n, D_n] : \mu_n \in \mathbb{R}^n \text{ \& } D_n \text{ is any } n \times n \text{ p.d. diagonal matrix} \right\}$$

where $g[\mu_n, D_n]$ is a normal density with mean μ_n and diagonal covariance $\sigma_f^2 D_n$. We represent the minimal Gaussian predictive risk by $\rho_{\mathcal{G}}(\theta_n) := \inf_{\hat{p} \in \mathcal{G}_n} \rho(\theta_n, \hat{p})$.

Our shrinkage results will mostly refer to the sub-family $\mathcal{G}_n[1]$ of $\mathcal{G}_n[p]$. $\mathcal{G}_n[1]$ contains Gaussian densities with only one data-adaptive scale estimate

$$\mathcal{G}_n[1] = \left\{ g[\mu_n, c] : \mu_n \in \mathbb{R}^n \text{ and } c \in \mathbb{R}^+ \right\}.$$

where $g[\hat{\theta}_n, c]$ denotes a normal density with mean μ_n and covariance $c \sigma_f^2 I$. A typical density estimate in $\mathcal{G}_n[1]$ is represented as $g[\hat{\theta}_n, \hat{c}(n)]$ where $\hat{\theta}_n$ is a location estimate and $\hat{c}(n)$ is the scale estimate based on observing an n -dimensional past observation X_n . For any fixed location estimate $\hat{\theta}_n$, the optimal risk of density estimates in $\mathcal{G}_n[1]$ centered around $\hat{\theta}_n$ is given by

$$\rho_0(\theta_n, \hat{\theta}_n) = \inf_{\hat{c}(X_n) \in \mathbb{R}^+} \rho(\theta_n, g[\hat{\theta}_n, \hat{c}(X_n)]).$$

The quadratic risk of the location estimate is denoted by

$$q(\theta_n, \hat{\theta}_n) = \mathbb{E}_{\theta_n} \|\hat{\theta}(X_n) - \theta_n\|^2$$

where the expectation is over the observed past X_n . Later, we show that if the value of $q(\theta_n, \hat{\theta}_n)$ were known, then the optimal choice for scale is

$$\text{IF}_{\theta_n}(\hat{\theta}_n) = 1 + n^{-1} r^{-1} q(\theta_n, \hat{\theta}_n)$$

which will be called as the Ideal Flattening coefficient for $\hat{\theta}_n$ at θ_n . Here, given a location estimate $\hat{\theta}_n$ we construct suitable estimates $\hat{c}(n)$ of the scale such that asymptotically when $n \rightarrow \infty$ the density estimate $g[\hat{\theta}_n, \hat{c}(n)]$ is optimally flattened in the sense that $\rho(\theta_n, g[\hat{\theta}_n, \hat{c}(n)]) - \rho_0(\theta_n, \theta_n) \leq O(1)$. However, for proving optimality of the flattening coefficient we need the following mild regularity conditions on the location estimate $\hat{\theta}_n$:

$$(3) \quad q(\theta_n, \hat{\theta}_n) \leq O(n).$$

$$(4) \quad \text{Var}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2) \leq O(n)$$

and the existence of a suitable estimate $U[\hat{\theta}](X_n)$ for the quadratic risk of $\hat{\theta}_n$ at θ_n with the following properties:

$$(5) \quad \left| \mathbb{E}_{\theta_n}(\hat{U}_n) - q(\theta_n, \hat{\theta}_n) \right| \leq O(n^{1/2}).$$

$$(6) \quad \text{Var}_{\theta_n}(\hat{U}_n) \leq O(n).$$

$$(7) \quad \text{Var}_{\theta_n} \left\{ \left(1 + (nr)^{-1} \hat{U}_n \right)^{-1} \right\} \leq O(n^{-1}).$$

These properties are fairly mild and in Section 2 we show that most popular point estimators obey these above conditions. We call these conditions Reasonable Asymptotic Square Loss (RASL) properties and the set of point estimators in the sequence model (where the action set is \mathbb{R}^∞) which satisfies these conditions is denoted by \mathcal{A} . Also, we denote the ratio of the future to past variances by $r := \sigma_f^2 / \sigma_p^2$. Our results will depend on r . For sequences, the symbol $a_n \sim b_n$ means $a_n = b_n(1 + o(1))$ and $a_n \approx b_n$ means $a_n/b_n \in (k_1, k_2)$ where k_1 and k_2 are constants.

Results. We show that in high dimensions, the minimum predictive entropy risk of Gaussian density estimates around reasonable location estimate $\hat{\theta}_n$ can be expressed in terms of the corresponding quadratic risk of $\hat{\theta}_n$. The minimum predictive risk can be attained by optimally flattening the normal density estimate around $\hat{\theta}_n$. The choice of the optimal flattening coefficient is not unique. An asymptotically efficient choice based on a reasonable estimate $U[\hat{\theta}_n](X_n)$ of the quadratic risk of $\hat{\theta}_n$ can be made.

THEOREM 1.1. *For any estimator $\hat{\theta}$ in \mathcal{A} we have*

$$(8) \quad \left| \rho_0(\theta_n, \hat{\theta}_n) - \frac{n}{2} \log \left(1 + (nr)^{-1} \cdot q(\theta_n, \hat{\theta}_n) \right) \right| \leq O(1) \quad \text{as } n \rightarrow \infty.$$

And if $\hat{c}(X_n) = 1 + (nr)^{-1} U[\hat{\theta}](X_n)$ is based on a suitable estimate \hat{U}_n of the quadratic risk as defined in Equations (5)–(7) then

$$(9) \quad \rho(\theta_n, g[\hat{\theta}_n, \hat{c}(n)]) - \rho_0(\theta_n, \hat{\theta}_n) \leq O(1).$$

We represent the optimal density estimate $g[\widehat{\theta}_n, \widehat{c}(n)]$ by $g[\widehat{\theta}_n]$. Based on the asymptotic relations between $\rho_0(\theta_n, \cdot)$ and the Mean Square Error (MSE) $q(\theta_n, \cdot)$, we can characterize the predictive risk of $g[\widehat{\theta}_n]$ easily by plugging in standard oracle inequalities from point estimation theory. We check that RASL conditions defined in Equations (3)–(7) hold for the James-Stein estimator (Stein, 1981)

$$\widehat{\theta}_n^{JS} = X_n \left(1 - \frac{n-2}{\|X_n\|^2} \right)$$

and its positive part estimator $\widehat{\theta}_n^{JS+}$. For the James-Stein estimator we determine the deviations from the optimal risk in terms of dimension dependent bounds.

THEOREM 1.2. *For any dimension $n \geq 10$ and for any $\theta_n \in \mathbb{R}^n$ we have,*

$$\rho(\theta_n, g[\widehat{\theta}_n^{JS}]) - \rho_0(\theta_n, \widehat{\theta}_n^{JS}) \leq 2^{-1} (a_n^{1/2} b_n^{1/2} r^{-3/2} + (a_n + b_n + l_n) r^{-2} + a_n r^{-3})$$

where the constants a_n, b_n, l_n are independent of the parameter but depend on the dimensions n and are given by

$$\begin{aligned} a_n &= 3(1 - (n-2)^{-1})^{-2}, \quad b_n = 4(2 + a_n + k_2(n)), \\ l_n &= 3(1 - 2/n)^{-2}, \quad k_2(n) = \max\{e(n), f(n)\} \text{ with} \\ e_n &= \sqrt{3} \prod_{i=1}^4 (1 - (2i+1)/n)^{-1/2} \text{ and } f_n = (1 - (\log n/n)^{1/2})^{-2}. \end{aligned}$$

Also, $\rho(\theta_n, g[\widehat{\theta}_n^{JS}])$ can be approximated by using the following bound

$$\left| \rho(\theta_n, g[\widehat{\theta}_n^{JS}]) - \frac{n \log IF_{\theta_n}(\widehat{\theta}_n)}{2} \right| \leq \frac{(a_n^{1/2} b_n^{1/2} r^{-3/2} + (a_n + b_n) r^{-2} + a_n r^{-3})}{2}.$$

These bounds hold for any value of $r \in (0, \infty)$. As the ratio of the future to past variances r decreases, we need to estimate the future observations based on increasingly noisy past observations and so, the difficulty of the density estimation problem also increases. So, as expected when r decreases the bounds also increases. These bounds can be made dimension independent. In particular, for all dimension $n \geq 20$, for any $\theta_n \in \mathbb{R}^n$ and for any fixed value of $r \in (0, \infty)$, we have,

$$(10) \quad \rho(\theta_n, g[\widehat{\theta}_n^{JS}]) - \rho(\theta_n, \widehat{\theta}_n^{JS}) \leq 5.3 r^{-3/2} + 19.6 r^{-2} + 1.7 r^{-3}.$$

In point estimation theory, there exist sharp oracle bounds on the quadratic risk $q(\theta_n, \hat{\theta}_n^{JS})$ of the James-Stein estimator $\hat{\theta}^{JS}$ (Johnstone, 2012, Chapter 2) which along with Theorem 1.2 produce the following oracle bound on the predictive risk of shrinkage predictive density estimates. Assuming that the value $\|\theta_n\|^2$ is known, the risk of the ideal linear predictive density estimate is given by

$$(11) \quad \text{IL}(\theta_n) = \frac{n}{2} \log \left(1 + r^{-1} \frac{a_n}{1 + a_n} \right) \text{ where } a_n = \|\theta_n\|^2/n.$$

The difference in the risk of $g[\hat{\theta}_n^{JS}]$ and the optimal oracle linear risk is

$$(12) \quad \rho(\theta_n, g[\hat{\theta}_n^{JS}]) - \text{IL}(\theta_n) \leq 0.1 r^{-1} + 5.3 r^{-3/2} + 18.1 r^{-2} + 1.7 r^{-3}.$$

Comparing this to the oracle bound of Xu and Zhou (2011) which is derived based on an empirical Bayes perspective

$$(13) \quad \rho(\theta_n, g[\hat{\theta}_n^{JS}]) - \text{IL}(\theta_n) \leq 2 r^{-1} + 5 r^{-2} + 4 r^{-3},$$

the particular features of our moment based approach can be seen. As our oracle inequality is a by-product of the optimal Gaussian risk, for most values of r the bound in the Inequality (12) is coarser than that in Inequality (13). However, when $r = 0.1$, the RHS in the Inequality (12) is 3830 and is better than the bound (4520) in the latter. Thus, the moment based approach can be quite informative. The bounds derived on the predictive risk are sharp enough to derive decision-theoretic optimality. We can produce unrestricted improved minimax predictive densities which asymptotically behaves like ideally shrunk linear density estimates (as defined later). The following lemma shows the asymptotic improvement in the predictive risk over the best invariant predictive density $g[X_n, 1 + r]$ (Liang and Barron, 2004).

LEMMA 1.1. *If $\|\theta_n\|^2 \rightarrow \infty$ as $n \rightarrow \infty$, then we have*

- [a] $\rho(\theta_n, g[X_n, 1 + r]) = 2^{-1} n \log(1 + r^{-1})$.
- [b] $\rho(\theta_n, g[\hat{\theta}_n^{JS}, r]) \sim (2r)^{-1} n a_n (1 + a_n)^{-1}$ where $a_n = n^{-1} \|\theta_n\|^2$.
- [c] $\rho(\theta_n, g[\hat{\theta}_n^{JS}, 1 + r]) \sim 2^{-1} n \{ \log(1 + r^{-1}) - (1 + a_n)^{-1} (1 + r)^{-1} \}$.
- [d] $\rho(\theta_n, g[\hat{\theta}_n^{JS}]) \sim 2^{-1} n \log \{ 1 + r^{-1} a_n (1 + a_n)^{-1} \}$.

The improvement in predictive risk due to efficient choice of location is reflected by the risk of $g[\hat{\theta}^{JS}, 1 + r]$ where as the effect of the optimal choice of scale after choosing an appropriate location estimate can be followed by evaluating the asymptotic predictive risk of $g[\hat{\theta}^{JS}]$.

The regularity conditions that we impose on the location point estimates do not extend to convex collections of estimates in \mathcal{A} . But, the predictive risk still concentrates and the optimal predictive risk ρ_0 can be determined.

LEMMA 1.2. *For any countable collection Λ of estimators $\hat{\theta}[\lambda]$ in \mathcal{A} and their convex collection $\hat{\theta}^w = \sum_{\lambda \in \Lambda} w_\lambda \hat{\theta}[\lambda]$ with $\sum_{\lambda \in \Lambda} w_\lambda = 1$, we have*

$$\rho_0(\theta_n, \hat{\theta}_n^w) - \frac{n}{2} \sum_{\lambda \in \Lambda} w_\lambda \log \left(1 + (nr)^{-1} \cdot q(\theta_n, \hat{\theta}_n[\lambda]) \right) \leq O(1) \text{ as } n \rightarrow \infty.$$

And, the predictive density estimate $\sum_{\lambda \in \Lambda} w_\lambda g[\hat{\theta}_n[\lambda]]$ is asymptotically optimal in the sense

$$(14) \quad \rho \left(\theta_n, \sum_{\lambda \in \Lambda} w_\lambda g[\hat{\theta}_n[\lambda]] \right) - \rho_0(\theta_n, \hat{\theta}_n^w) \leq O(1) \text{ as } n \rightarrow \infty.$$

The class $\mathcal{G}[p]$ of product Gaussian density estimates is minimax optimal over ellipsoids (Xu and Liang, 2010). However, we show that the class $\mathcal{G}[p]$ is minimax sub-optimal over the ℓ_0 -sparsity constrained space $\Theta(n, s)$ when $s/n \rightarrow 0$ as $n \rightarrow \infty$.

THEOREM 1.3. *For any fixed $r \in (0, \infty]$ as $n \rightarrow \infty$ for every sequence s_n with $s_n/n \rightarrow 0$, we have*

$$(15) \quad \min_{\hat{p} \in \mathcal{G}_n[p]} \max_{\Theta(n, s)} \rho(\theta_n, \hat{p}) = r^{-1} s \log(n/s)(1 + o(1)).$$

By Theorem 1.2 in Mukherjee and Johnstone (2012), we know that the asymptotic minimax risk $R(n, s, r)$ over $\Theta(n, s)$ is given by

$$(16) \quad R(n, s, r) \sim (1 + r)^{-1} s \log(n/s) \text{ as } n \rightarrow \infty, s \rightarrow \infty \text{ and } s/n \rightarrow 0.$$

Hence, the minimax sub-optimality of the class $\mathcal{G}[p]$ over the ℓ_0 sparse space is $1 + r^{-1}$. The parametric space $\Theta(n, s)$ is not invariant to the group of orthogonal transformations. If the parameter space does not have any sparser representation with respect to the group of orthogonal transformations, then the asymptotic sub-optimality of the class \mathcal{G} is also $1 + r^{-1}$.

1.3. *Organization of the paper.* The predictive error of the class $\mathcal{G}[1]$ of predictive density estimates is presented in the next section, which discusses the role of shrinkage in high-dimensional prediction problems. Predictive error and restricted minimax risk of the class \mathcal{G} is presented in Section 3.

2. Role of shrinkage and optimal error in $\mathcal{G}[1]$.

Hereon we will assume that $\sigma_p^2 = 1$ and $\sigma_f^2 = r$. The general predictive KL risk will not be affected by this restriction. However, the density estimates are usually based on statistics equivariant to the scale transformation and needs multiplication by σ_p .

Heuristic Idea. In the high dimensions the quadratic loss of a reasonable point estimator will concentrate around its risk. And, so the KL risk of the corresponding Gaussian predictive density partitions into two parts involving (i) quadratic risk on the location parameter adjusted by the expected scale (ii) logarithm of the expected scale. As such, the risk $\rho(\theta_n, g[\hat{\theta}_n, \hat{c}_n])$ of the normal predictive density estimate $g[\hat{\theta}_n, \hat{c}_n]$ is given by

$$\frac{n}{2} \left[\left\{ \mathbb{E}_{\theta_n}(\log \hat{c}(X_n)) + \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}(X_n)} \right) - 1 \right\} + \frac{1}{nr} \mathbb{E}_{\theta_n} \left(\frac{\|\hat{\theta}(X_n) - \theta_n\|^2}{\hat{c}(X_n)} \right) \right].$$

In high dimensions, due to concentration of measure we expect

- $\mathbb{E}_{\theta_n} \log(\hat{c}_n) \sim \log \mathbb{E}_{\theta_n} \hat{c}_n$
- $\mathbb{E}_{\theta_n}(\hat{c}_n^{-1}) \sim (\mathbb{E}_{\theta_n} \hat{c}_n)^{-1}$
- $\mathbb{E}_{\theta_n} \left(\|\hat{\theta}(X_n) - \theta_n\|^2 \cdot \hat{c}_n^{-1}(X_n) \right) \sim (\mathbb{E}_{\theta_n} \hat{c}_n^{-1}(X_n))^{-1} q(\theta_n, \hat{\theta}_n)$

which will lead to

$$\rho(\theta_n, g[\hat{\theta}_n, \hat{c}_n]) \sim \frac{n}{2} \left\{ \log \mathbb{E}_{\theta_n} \hat{c}_n + \frac{1 + (nr)^{-1} q(\theta_n, \hat{\theta}_n)}{\mathbb{E}_{\theta_n} \hat{c}_n} - 1 \right\} + O(1) \text{ as } n \rightarrow \infty.$$

This asymptotic decomposition of the predictive risk can be explicitly validated through the RASL properties. Because of this decomposition for any fixed point estimate $\hat{\theta}_n$ at each parametric value θ_n we can minimize the above asymptotic value of $\rho(\theta_n, g[\hat{\theta}_n, \cdot])$ over the scalar quantity $\mathbb{E}_{\theta_n} \hat{c}_n$. The minimum asymptotic value is given by

$$\rho(\theta_n, g[\hat{\theta}_n]) \sim n/2 \cdot \log(1 + (nr)^{-1} q(\theta_n, \hat{\theta}_n))$$

and the optimal value is attained when

$$\mathbb{E}_{\theta_n}(\hat{c}^{\text{opt}}(X_n)) = 1 + (nr)^{-1} q(\theta_n, \hat{\theta}_n) = \text{IF}_{\theta_n}(\hat{\theta}_n)$$

which is the ideal flattening coefficient. However, $\text{IF}_{\theta_n}(\hat{\theta}_n)$ is unknown. But, it depends only on the parametric value θ_n . Thus a choice would be $\hat{c}^{\text{opt}}(X_n) =$

$1 + (nr)^{-1}U[\widehat{\theta}_n](X_n)$, where $U[\widehat{\theta}_n](X_n)$ (to be abbreviated as \widehat{U}_n) is reasonable (i.e with reasonable bias and concentration properties) estimate of the quadratic risk of $\widehat{\theta}_n$. With very high probability such an optimal choice of \widehat{c}_n will be greater than 1 reflecting a flattening of scale of the estimated predictive density (with respect to the true future variability). Intuitively, we are performing an appropriate flattening of the density based on empirical estimates of the quadratic loss. The optimal density $g[\widehat{\theta}_n]$ levels out with increasing inaccuracy in the location estimate $\widehat{\theta}_n$.

One of the popular Frequentist notion (which is better than plug-in density estimates) of constructing predictive densities in this parametric model is to use Gaussian density estimate around an efficient location $\widehat{\theta}_n$ and variance $r + \widehat{\text{Var}}(\widehat{\theta}_n)$. Estimates of these kind are natural extensions of confidence sets. The optimal density estimate $g[\widehat{\theta}]$ is quite similar except with a larger variance $r + \widehat{q}(\theta_n, \widehat{\theta}_n)$. And unless the bias of θ_n is negligible compared to its variance the above mention general notion produces sub-optimal density estimates. Next through the RASL conditions we will quantify some statistical regularities in the behavior of quadratic loss in high dimensions.

2.1. RASL Properties of a Location Point Estimate. In high dimensions, for any fixed location parameter θ_n and its estimate $\widehat{\theta}_n$ we expect the quadratic loss $\|\widehat{\theta}_n - \theta_n\|_2^2$ to be concentrated around its expected value (quadratic risk) $\mathbb{E}_{\theta_n}\|\widehat{\theta}_n - \theta_n\|_2^2$ and it would be reflected by its variance. We will also rule out very bad point estimators by neglecting those with too high risk as we do not want them for prediction purposes. Apart from these we also assume the existence of a statistic which estimates the quadratic risk within reasonable bias. These properties of point estimators are referred to as **R**easonable **A**symptotic **S**quare **L**oss properties and the corresponding location estimates as RASL estimates.

As dimension $n \rightarrow \infty$, for any fixed parametric value θ_n the location point estimate $\widehat{\theta}(X_n)$ is such that its quadratic loss has the following properties.

P1. Reasonable Risk:

$$\mathbb{E}_{\theta_n}\|\widehat{\theta}_n - \theta_n\|^2 \leq \mathbf{O}(n).$$

The canonical minimax point estimator X_n which is also the UMVUE (under square loss) in this case acts as the benchmark in weeding out the bad point estimators. For any parameter value θ_n , X_n has constant risk n . So, it is appropriate for our purpose to restrict ourselves to point estimators with risk of the $\mathbf{O}(n)$.

P2. Concentration property of Quadratic loss:

$$\mathbb{V}\text{ar}_{\theta_n} (\|\hat{\theta}_n - \theta_n\|^2) \leq \mathbf{O}(\mathbf{n}).$$

In high dimensions the estimator $\hat{\theta}_n$ is such that its loss has variability less than $O(n)$. Again comparing with X_n , we see $\mathbb{V}\text{ar}_{\theta_n} (\|X_n - \theta_n\|^2) = 2n$ as $\|X_n - \theta_n\|^2$ is distributed as a central χ^2 random variable with n degrees of freedom.

P2 implies concentration of the loss function and would in turn also impose some concentration properties on well-behaved functions of the loss. As such, using Lemma **A.1**, we have

$$\mathbf{P2.a} \quad \mathbb{V}\text{ar}_{\theta_n} \left\{ \left(1 + \frac{\|\hat{\theta}_n - \theta_n\|^2}{\mathbf{n}\mathbf{r}} \right)^{-1} \right\} \leq \mathbf{O}(\mathbf{n}^{-1})$$

following directly from **P2**. It is an important condition and will be used in our derivations.

P3. Reasonable Estimate of Quadratic Risk:

There exists an estimator $U[\hat{\theta}_n](X_n)$ (will be abbreviated as \hat{U}_n) of the quadratic risk of $\hat{\theta}_n$ satisfying the following:

$$\mathbf{P3.1.} \quad \left| \mathbb{E}_{\theta_n}(\hat{U}_n) - \mathbb{E}_{\theta_n} \|\hat{\theta}_n - \theta_n\|^2 \right| \leq \mathbf{O}(\mathbf{n}^{1/2}).$$

$$\mathbf{P3.2.} \quad \mathbb{V}\text{ar}_{\theta_n}(\hat{U}_n) \leq \mathbf{O}(\mathbf{n}).$$

$$\mathbf{P3.3.} \quad \mathbb{V}\text{ar}_{\theta_n} \left\{ \left(1 + \frac{\hat{U}_n}{\mathbf{n}\mathbf{r}} \right)^{-1} \right\} \leq \mathbf{O}(\mathbf{n}^{-1}).$$

P3.1 implies existence of a statistic which estimates the quadratic risk by not making significant bias. Bias exceeding $O(\sqrt{n})$ is considered significant here and the order is associated with the $O(n^{-1})$ asymptotic statements we would like to make. **P3.1** and **P3.2** are analogous to **P2** and **P2.a** respectively. They imply that the asymptotic concentration properties associated with the quadratic loss also holds for its estimator \hat{U}_n . If \hat{U}_n is positive then **P3.2** follows directly from **P3.1** by Lemma **A.1**.

2.2. *Validating the RASL properties.* Given a location point estimator and its corresponding reasonable quadratic risk estimate the RASL conditions can be checked at least by simulations. However, existence of a ‘reasonable’ risk estimator (as defined in **P3**) is essential. For most widely used point estimates, we can construct risk estimates satisfying the three conditions in **P3** though the procedures can sometime get quite complicated.

If $\hat{\theta}_n$ is the posterior mean – generalized Bayes estimate with respect to prior π , then by Tweedie’s formula [Brown \(1971\)](#), [Robbins \(1956\)](#) we have explicit expression of an unbiased estimate of its risk as,

$$\hat{U}_n^\pi = n - \left[\|\nabla \log m_\pi(X_n)\|^2 - 2 \frac{\nabla^2 m_\pi(X_n)}{m_\pi(X_n)} \right] \text{ where } \nabla f \triangleq \sum_{i=1}^n D_i f$$

$$\nabla^2 m_\pi(X_n) = \sum_{i=1}^n D_i^2 m_\pi(X_n) \text{ and } m_\pi(x_n) = \int \phi_n(x_n|\theta_n, 1) \pi(\theta_n) d\theta_n.$$

\hat{U}_n^π is a natural candidate for a ‘reasonable estimate of the quadratic loss’ though **P3.2** and **P3.3** are also to be checked separately. In particular, for **P3.3** to hold \hat{U}_n^π may need some modification by introducing some bias.

For spherically symmetric estimators, we can get candidates for ‘reasonable’ risk estimates by using Stein’s unbiased (quadratic) risk estimates (SURE) or their modifications (like positive part, etc) [Stein \(1974, 1981\)](#). As mentioned before, here too we needed to introduce some bias to the the SURE estimate as the unbiased one does not has property P3.3.

These RASL conditions are quite mild and usually holds for reasonable point estimates and can be checked by Monte Carlo simulations for arbitrary point estimates. Next, we check these conditions analytically for the following popular point estimators:

$\hat{\theta}^{JS}$: James Stein estimator

$\hat{\theta}^{JS+}$: Positive part James Stein estimator

$\hat{\theta}^H$: Posterior mean of harmonic prior $\pi_H(\theta_n) \propto \|\theta_n\|^{-(n-2)}$.

All these 3 point estimators are linear estimates of the form $s(X_n)X_n$ where $s(X_n)$ is a data-dependent shrinkage term. They are better than the canonical minimax estimator X_n . While $\hat{\theta}^H$ is admissible, $\hat{\theta}^{JS}$ and $\hat{\theta}^{JS+}$ are both inadmissible. As such both $\hat{\theta}^{JS+}$ and $\hat{\theta}^H$ dominates $\hat{\theta}^{JS}$. However, in high dimensions, they behave similarly and have near ideal linear risk properties. We will construct reasonable risk estimates for each of these estimators. While verifying the RASL conditions for the JS estimator we would also compute the bound explicitly for each n . It will be needed afterwards

in Theorem 1.2. Since, the estimators are spherically symmetric it will be more informative to derive bounds depending on $\|\theta_n\|^2$. Hence forth in this section, by a_n we denote $\|\theta_n\|^2/n$. A convenient fact about this spherically symmetric estimators is that the n -dimensional parameter θ_n can be substituted by $(\|\theta_n\|, 0, \dots, 0)$ while checking the asymptotic behavior of square loss. As these estimators are not Lipchitz functions of the normal random variable X , we can not directly use well-established Gaussian concentration inequalities (Dembo and Zeitouni, 1993, Ledoux, 2001).

2.2.1. *James Stein estimator.* The James-Stein estimator and its unbiased risk estimate is given by:

$$\hat{\theta}_n^{JS} = X_n \left(1 - \frac{n-2}{\|X_n\|^2} \right), \quad \text{and} \quad U(\hat{\theta}_n^{JS}) = \left(n - \frac{(n-2)^2}{\|X_n\|^2} \right).$$

RASL property **P1.** holds as the JS is better than the canonical estimator X_n . As such a good upper bound on its risk is also known

$$E_{\theta_n} \|\hat{\theta}_n^{JS} - \theta_n\|^2 \leq 2 + \frac{(1 - 2/n) a_n}{(1 - 2/n) + a_n}.$$

LEMMA 2.1.

$$\mathbb{V}ar_{\theta_n} \left(\|\hat{\theta}_n^{JS} - \theta_n\|^2 \right) \leq 4 \left[2n + (n-2)^4 n^{-3} k_1(n) + n k_2(n) \right].$$

PROOF. We decompose $\|\hat{\theta}_n^{JS} - \theta_n\|^2$ into 3 parts as

$$\|\hat{\theta}_n^{JS} - \theta_n\|^2 = \|X_n - \theta_n\|^2 + (n-2)^2 \|X_n\|^{-2} + 2(n-2)M_n$$

where $M_n = \langle X_n - \theta_n, X_n \|X_n\|^{-2} \rangle$. Then we use the naive inequality that for any three random variables $Z_i, i = 1, 2, 3$

$$\mathbb{V}ar \left(\sum_{i=1}^3 Z_i \right) \leq \sum_{j=0}^3 \mathbb{V}ar \left(\sum_{i=1}^3 (-1)^{\mathbb{I}\{j=i\}} Z_i \right) = 4 \sum_{i=1}^3 \mathbb{V}ar(Z_i)$$

to get the following bound on $\mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2)$

$$\leq 4 \left\{ \mathbb{V}ar_{\theta_n}(\|X_n - \theta_n\|^2) + \mathbb{V}ar_{\theta_n} \left(\frac{n-2}{\|X_n\|^2} \right) + 4(n-2)^4 \mathbb{V}ar_{\theta_n}(M_n) \right\}.$$

Now $\|X_n - \theta_n\|^2$ has a central chi-square distribution with n degrees of freedom and hence its variance is $2n$. The bounds on the other quantities follow from Lemma 2.2 and Lemma A.1. \square

LEMMA 2.2. For $n \geq 10$ we have $\mathbb{V}ar_{\theta_n}(M_n) \leq n^{-1}k_2(n)$ where

$$k_2(n) = \max\{h(n), k(n)\} \text{ where } e_n = \frac{\sqrt{3}}{\prod_{i=1}^4 (1 - (2i+1)/n)}^{1/2}$$

$$\text{and } f_n = \frac{1}{(1 - (\log n/n)^{1/2})^2}.$$

PROOF. The variance of M_n is same as the variance of $\langle \theta_n, X_n \rangle \|X_n\|^{-2}$ whose distribution is spherically symmetric in θ_n as it can be written as sum of two spherically symmetric terms $H_n = \langle \theta_n, X_n - \theta_n \rangle \|X_n\|^{-2}$ and $J_n = \|\theta_n\|^2 \|X_n\|^{-2}$. So, with out loss of generality we can assume that $\theta_n = (\theta, 0, \dots, 0)$ where $\theta = \|\theta_n\|$. We also divide the proof into two cases depending on the magnitude of θ .

When $\theta \leq \sqrt{n}$ we have, $\mathbb{V}ar_{\theta_n}(M_n) \leq 2(\mathbb{V}ar_{\theta_n}(H_n) + \mathbb{V}ar_{\theta_n}(J_n))$ with the later being less than n^{-1} by Lemma A.3. And, the former is bounded above by $E(H_n^2)$. Now, with $Z \stackrel{d}{=} N(0, 1)$ and $W \stackrel{d}{=} \chi_{n-1}^2(0)$ and $V = (Z + \theta)^2 + W$ it can be rewritten as

$$E(\theta^2 Z^2 V^{-2}) \leq \sqrt{\theta^4 E Z^4 E V^{-4}} \leq \sqrt{3\theta^4 E(W^{-4})} \leq \left\{ \frac{3\theta^4}{\prod_{i=1}^4 (n - 2i - 1)} \right\}^{1/2}$$

which is less than $n^{-1}\sqrt{3}\prod_{i=1}^4 (1 - (2i+1)/n)^{-1/2}$.

When $\theta > n$, we first recall that $M_n \stackrel{d}{=} (\theta + Z)/W$ and so

$$\begin{aligned} E(M_n^2) &\leq E\{V^{-2}\mathbb{I}_{\{|\theta+Z|\leq 1\}}\} + E\{(\theta + Z)^{-2}\mathbb{I}_{\{|\theta+Z|>1\}}\} \\ &\leq E\{V^{-2}\} + 2 \int_1^\infty x^{-2}\phi(x - \theta) dx \\ &\leq [(n-3)(n-5)]^{-1} + \tilde{\Phi}(\sqrt{\log n}) + \{\theta - \sqrt{\log n}\}^{-2} \\ &\leq [(n-3)(n-5)]^{-1} + n^{-1}(\log n)^{-1} + n^{-1}(1 - (\log n/n)^{1/2})^{-2} \end{aligned}$$

Hence the result follows. \square

Though it is very tempting but we can not use the unbiased risk estimate $U(\hat{\theta}_n^{JS})$ as the estimate can be negative and violates **P3.3**.

LEMMA 2.3. For any fixed n and r , $E_0[\{1 + (nr)^{-1}U(\hat{\theta}_n^{JS})\}^{-1}]$ does not exist.

We will instead use \hat{U}_n^+ the positive part of $U(\hat{\theta}_n^{JS})$ and the scale estimate $\hat{c}_n^{JS+} = 1 + (nr)^{-1}\hat{U}_n^+$. RASL condition **P3.1** can be easily checked as

$$\mathbb{V}ar_{\theta_n}(\hat{U}_n^+) \leq \mathbb{V}ar_{\theta_n}(U(\hat{\theta}_n^{JS})) = (n-2)^4 \mathbb{V}ar_{\theta_n}(\|X_n\|^{-2}) = O(n)$$

by Lemma A.1 and P3.2 follows from Lemma A.3. As such, an exact dimension dependent bound can also be derived.

LEMMA 2.4. *For any fixed $n \geq 3$ we have*

$$|\text{Bias}_{\theta_n}(\hat{c}_n^{JS+})| \leq k_3(n) n^{-1/2} \text{ where } k_3(n) = \frac{\sqrt{2} + 5n^{-1/2}}{1 - 2/n}.$$

PROOF. Noting that $\text{Bias}_{\theta_n}(\hat{c}_n^{JS+}) = n^{-1/2} \mathbb{E}_{\theta_n}(\hat{U}_n^-)$ and

$$n |\mathbb{E}_{\theta_n}(\hat{U}_n^-)| \leq \mathbb{E}_{\theta_n} \left[\left(\frac{n}{Y} - 1 \right) \cdot \mathbb{I}\{Y \leq n\} \right]$$

where Y follows Chi-square with degree n and non-centrality parameter $\|\theta_n\|^2$. We know that $Y \stackrel{d}{=} \chi_{n+2N}^2$ where $N \stackrel{d}{=} \text{Poisson}(\|\theta_n\|^2/2)$ and the above expectation can be written as

$$\mathbb{E}_{\|\theta_n\|^2} \left(\mathbb{E} \left[\left(\frac{n}{Y_{n+2N}} - 1 \right) \cdot \mathbb{I}\{Y_{n+2N} \leq n\} \middle| N \right] \right) \leq \mathbb{E} \left[\left(\frac{n}{Y_n} - 1 \right) \cdot \mathbb{I}\{Y_n \leq n\} \right]$$

where Y_{n+2N} is a central chi-square random variable with $(n+2N)$ degrees of freedom and the second inequality follows as for any $N \geq 0$, $(n/Y_{n+2N} - 1) \cdot \mathbb{I}\{Y_{n+2N} \leq n\}$ is stochastically dominated by $N = 0$. Now,

$$\mathbb{E} \left[\left(\frac{n}{Y} - 1 \right) \cdot \mathbb{I}\{Y \leq n\} N \right] = \int_0^n \frac{n-y}{y} \frac{y^{n/2-1} e^{-y/2}}{2^{n/2} \Gamma(n/2)} dy \leq \frac{E|W_n - n|}{n-2}$$

where $W_n \sim \text{Gamma}(n/2 - 1, 1/2)$ and so

$$E|W_n - n| \leq 1 + E(W_n - n)^2 \leq 1 + 4 + \text{Var}(W_n) = 5 + \sqrt{2n}$$

where the second inequality follows by Bias-Variance decomposition. Thus, we get our result. \square

2.2.2. *James-Stein Positive part Estimator.* We consider the positive part of the JS estimator and a reasonable estimate of its loss as

$$\hat{\theta}_n^{JS+} = X_n \left(1 - \frac{n-2}{\|X_n\|^2} \right)_+ \quad \text{and} \quad U(\hat{\theta}_n^{JS+}) = \left(n - \frac{(n-2)^2}{\|X_n\|^2} \right)_+.$$

There exists unbiased estimator of the quadratic risk of $\hat{\theta}_n^{JS+}$ (Johnstone, 2012, Exercise 2.13). We use a biased estimator here mainly to highlight the fact that even biased estimators will work.

P1. follows from the fact that $\hat{\theta}_n^{JS+}$ is better than $\hat{\theta}^{JS_n}$ (Johnstone, 2012, Exercise 2.8).

For checking **P2**, define C_n to be the event $\{\hat{\theta}_n^{JS+}(X_n) \neq 0\} = \{X_n : \hat{\theta}_n^{JS+} = \hat{\theta}_n^{JS}\}$. And the idea is to relate the variance of the loss in JS+ case with the case of JS estimator.

$$\begin{aligned}
\mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS+} - \theta_n\|^2) &= \mathbb{E}\|\hat{\theta}_n^{JS+} - \theta_n\|^4 - \mathbb{E}^2\|\hat{\theta}_n^{JS+} - \theta_n\|^2 \\
&= \mathbb{E}\{\|\hat{\theta}_n^{JS} - \theta_n\|^4 I_{C_n}\} - \mathbb{E}^2\{\|\hat{\theta}_n^{JS} - \theta_n\|^2 I_{C_n}\} + \|\theta_n\|^4 P(C_n^c) - \|\theta_n\|^4 P^2(C_n^c) \\
&= \mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2 | C_n) \cdot P_{\theta_n}(C_n) + \|\theta_n\|^4 P_{\theta_n}(C_n) P_{\theta_n}(C_n^c) \\
&\leq \mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2) + \|\theta_n\|^4 P_{\theta_n}(C_n^c) \\
&\text{as } \mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2) \geq \mathbb{E}_{\theta_n}(\mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2 | C_n))
\end{aligned}$$

We know that $\mathbb{V}ar_{\theta_n}(\|\hat{\theta}_n^{JS} - \theta_n\|^2)$ is $O(n)$ and lemma A.5 shows $\|\theta_n\|^4 P_{\theta_n}(C_n^c) \leq O(n)$. So, we have the desired bound.

Condition P.3.1. We will condition on the event C_n again and express **P.3.1** in terms of the James-Stein estimator

$$\mathbb{E}_{\theta_n} U(\hat{\theta}_n^{JS+}) - q(\theta_n, \hat{\theta}_n^{JS+}) = \mathbb{E}_{\theta_n} \left\{ \left(U(\hat{\theta}_n^{JS}) - \|\hat{\theta}_n^{JS} - \theta_n\|^2 \right) I_{C_n} \right\} - \|\theta_n\|^2 P(C_n^c).$$

When $\theta_n = 0$ then the R.H.S for large n reduces to,

$$I_n = E \left\{ \left(n - \frac{n^2}{Y_n} - \left(1 - \frac{n}{Y_n} \right)^2 Y_n \right) I_{C_n} \right\}$$

where Y_n is an central chi-squared random variable with n degrees of freedom. Now, we decompose I_n into

$$\begin{aligned}
I_n &= I_n^1 + 2I_n^2 \text{ where} \\
I_n^1 &= E \{(n - Y_n) I_{C_n}\} \text{ and } I_n^2 = E \{(n - n^2/Y_n) I_{C_n}\}
\end{aligned}$$

We standardize Y_n as $Z_n = (Y_n - n)/\sqrt{2n}$. We can use concentration inequalities on Z_n and have, [need to make rigorous]

$$\begin{aligned}
I_n^1 &\leq \sqrt{2n} \cdot EZ_+ \text{ as } n \rightarrow \infty \\
I_n^2 &= \sqrt{n} \cdot E \left\{ \left(\frac{Z_n}{Z_n/\sqrt{n} + 1} \right) I_{C_n} \right\} \leq \sqrt{n} E|Z_n| \rightarrow \sqrt{n} E|Z|
\end{aligned}$$

as on C_n , $Z_n \geq 0$. Thus $I_n \leq O(\sqrt{n})$.

Condition P3.2. By Lemma A.8 we have

$$\mathbb{V}ar_{\theta_n}(U(\hat{\theta}_n^{JS+})) \leq \mathbb{V}ar_{\theta_n}(U(\hat{\theta}_n^{JS})) = O(n).$$

Condition P3.3. Follows from Lemma A.1

Harmonic Prior The conditions can be checked for $\hat{\theta}^H$ by using its closed form expressions in (Xu, 2007, Chapter 2).

2.3. Determining ρ_0 for RASL point estimators. In this section, we will show that in high dimension with very high precision we can express $\rho_0(\hat{\theta}_n)$ – the minimum Predictive Entropy risk of the class of Gaussian density estimates around location $\hat{\theta}_n$ in terms of the Mean Square estimation error of θ_n by $\hat{\theta}_n$. We initially prove bounds on the error rates which holds for all dimensions but are dimension dependent. Then, we would show that in high dimensions those bounds are asymptotically sharp.

2.3.1. Lower Bound on $\rho_0(\theta_n, \hat{\theta}_n)$: . Next, we produce a lower bound on the prediction error. The bound ultimately will be a function of θ_n though it depends on the form of $\hat{\theta}_n$. It involves expectation of a quantity which usually is neither a parameter nor a statistic and hence can not be computed in closed form.

LEMMA 2.5. *For any dimension n , any parameter value θ_n and any location point estimate $\hat{\theta}(X_n)$, we have*

$$\rho_0(\theta_n, \hat{\theta}_n) \geq \frac{1}{2} \mathbb{E}_{\theta_n} \left\{ \log \left(1 + r^{-1} \cdot n^{-1} \cdot \|\hat{\theta}_n - \theta_n\|^2 \right) \right\}.$$

PROOF. For any fixed n , the risk of the predictive density q_n which is a n -dimensional product normal with with data adaptive mean $\hat{\theta}(X_n)$ and data and parameter dependent variance $\hat{c}(\theta_n, X_n) r$ (for all co-ordinates) is given by $2\rho(\theta_n, q_n)$

$$\begin{aligned} &= n \left\{ \mathbb{E}_{\theta_n}(\log \hat{c}(\theta_n, X_n)) + \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}(\theta_n, X_n)} \right) - 1 \right\} + \frac{1}{r} \mathbb{E}_{\theta_n} \left(\frac{\|\hat{\theta}(X_n) - \theta_n\|^2}{\hat{c}(\theta_n, X_n)} \right) \\ &= n \mathbb{E}_{\theta_n} \left\{ \log \hat{c}(\theta_n, X_n) + \frac{1 + (nr)^{-1} \|\hat{\theta}(X_n) - \theta_n\|^2}{\hat{c}(\theta_n, X_n)} - 1 \right\}. \end{aligned}$$

For any fixed value of θ_n and for each x_n ,

$$\log \hat{c}(\theta_n, x_n) + \hat{c}^{-1}(\theta_n, x_n) \{1 + (nr)^{-1} \|\hat{\theta}(x_n) - \theta_n\|^2\} - 1$$

is minimized at $\hat{c}^{\text{opt}}(\theta_n, x_n) = 1 + (nr)^{-1} \|\hat{\theta}(x_n) - \theta_n\|^2$ and the minimum value is given by $\log(1 + (nr)^{-1} \|\hat{\theta}(x_n) - \theta_n\|^2)$. Hence, the result follows. \square

Though $\hat{c}^{\text{opt}}(\theta_n, X_n)$ is the best possible flattening coefficient, it depends on the parameter and can not be used in practice. As such, $\hat{c}^{\text{opt}}(\theta_n, X_n)$ is the ideal flattening coefficient. In high dimensions due to statistical regularity we expect $\hat{c}^{\text{opt}}(\theta_n, X_n)$ to be very close to its expected value

$$\mathbb{E}_{\theta_n} \{\hat{c}^{\text{opt}}(\theta_n, X_n)\} = 1 + (nr)^{-1} \mathbb{E}_{\theta_n} \|\hat{\theta}_n - \theta_n\|^2$$

which can be viewed as the (near) **I**deal **F**lattening coefficient and is referred to as $\text{IF}_{\theta_n}(\hat{\theta}_n) = 1 + n^{-1}r^{-1}q(\theta_n, \hat{\theta}_n)$. Here flattening coefficients are usually called scale and it should be noted that the corresponding variance needs to be multiplied by r .

From Lemma 2.5 we can derive a worse but more tractable bound

$$(17) \quad \rho_0(\theta_n, \hat{\theta}_n) \geq 2^{-1} \mathbb{E}_{\theta_n} \left\{ \log \left(\|\hat{\theta}_n - \theta_n\|^2 / (nr) \right) \right\}.$$

2.3.2. Upper Bound for $\rho_0(\theta_n, \hat{\theta}_n)$: We now produce an upper bound on the risk of any Gaussian density estimate. Henceforth, SD would mean Standard Deviation and by Bias of the scale estimate \hat{c}_n we would mean the expected deviation from the near ideal flattening coefficient $\text{IF}_{\theta_n}(\hat{\theta}_n)$. With scales estimators based on the statistic $U[\hat{\theta}_n](X_n)$ and of the form $\hat{c}(X_n) = (1 + n^{-1}U[\hat{\theta}_n](X_n))$ we have $\text{Bias}_{\theta_n}(c_n) = (nr)^{-1} [\mathbb{E}_{\theta_n} U[\hat{\theta}_n](X_n) - q(\theta_n, \hat{\theta}_n)]$.

LEMMA 2.6. *For any fixed dimension n , parameter value θ_n , location point estimate $\hat{\theta}(X_n)$ and any scale estimate $\hat{c}(X_n) > 0$ almost surely and of the form $\hat{c}(X_n) = 1 + (nr)^{-1}U[\hat{\theta}_n](X_n)$, we have*

$$\rho(\theta_n, g[\hat{\theta}_n, \hat{c}_n]) - \frac{n}{2} \cdot \log(\text{IF}_{\theta_n}(\hat{\theta}_n)) \leq \frac{n}{2} \cdot \left[A_{\theta_n}(\hat{\theta}_n, \hat{c}_n) + B_{\theta_n}(\hat{\theta}_n, \hat{c}_n) \right]$$

$$\text{where } A_{\theta_n}(\hat{\theta}_n, \hat{c}_n) = \text{IF}_{\theta_n}(\hat{\theta}_n) \{ \mathbb{E}_{\theta_n}(\hat{c}_n) \}^{-1} \text{SD}_{\theta_n}(\hat{c}_n) \text{SD}_{\theta_n}(\hat{c}_n^{-1})$$

$$+ r^{-1} \text{SD}_{\theta_n} \left(\frac{\|\hat{\theta}_n - \theta_n\|^2}{n} \right) \text{SD}_{\theta_n}(\hat{c}_n^{-1}) \text{ and}$$

$$B_{\theta_n}(\hat{\theta}_n, \hat{c}_n) = \text{Bias}_{\theta_n}^2(\hat{c}_n) \{ \text{IF}_{\theta_n}(\hat{\theta}_n) \}^{-1} \{ \mathbb{E}_{\theta_n}(\hat{c}_n) \}^{-1}.$$

PROOF. The risk of the normal predictive density estimate $g[\hat{\theta}_n, \hat{c}_n]$ is given by $2\rho(\theta_n, \hat{g}_n)$

$$= n \left\{ \mathbb{E}_{\theta_n}(\log \hat{c}(X_n)) + \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}(X_n)} \right) - 1 \right\} + \mathbb{E}_{\theta_n} \left(\frac{\|\hat{\theta}(X_n) - \theta_n\|^2}{r \hat{c}(X_n)} \right).$$

Now, we replace $\mathbb{E}_{\theta_n}(\hat{c}_n^{-1})$ by $\mathbb{E}_{\theta_n}^{-1}\hat{c}_n$ and $\mathbb{E}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2 \times \hat{c}_n^{-1})$ by $\mathbb{E}_{\theta_n}\|\hat{\theta}_n - \theta_n\|^2 \times \mathbb{E}_{\theta_n}^{-1}\hat{c}_n$ in the above expression to get $\tilde{\rho}(\theta_n, \hat{g}_n)$

$$\begin{aligned}
(18) \quad &= \frac{1}{2} \left[n \left\{ \mathbb{E}_{\theta_n}(\log \hat{c}(X_n)) + \frac{1}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))} - 1 \right\} + \frac{1}{r} \left(\frac{\mathbb{E}_{\theta_n}\|\hat{\theta}(X_n) - \theta_n\|^2}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))} \right) \right] \\
&= \frac{n}{2} \mathbb{E}_{\theta_n}(\log \hat{c}(X_n)) + \frac{n}{2} \left\{ \frac{1 + (nr)^{-1} \mathbb{E}_{\theta_n}\|\hat{\theta}(X_n) - \theta_n\|^2}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))} - 1 \right\} \\
&= \frac{n}{2} \mathbb{E}_{\theta_n}(\log \hat{c}(X_n)) - \frac{n}{2} \frac{\text{Bias}_{\theta_n}(\hat{c}(X_n))}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))}
\end{aligned}$$

and the distortion caused thereby $(n/2)^{-1}(\rho(\theta_n, \hat{g}_n) - \tilde{\rho}(\theta_n, \hat{g}_n))$ equals

$$(19) \quad \mathbb{E}_{\theta_n} \left(\frac{1 + (nr)^{-1} \|\hat{\theta}(X_n) - \theta_n\|^2}{\hat{c}(X_n)} \right) - \frac{1 + (nr)^{-1} \mathbb{E}_{\theta_n} \|\hat{\theta}(X_n) - \theta_n\|^2}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))}.$$

Next we will show that $(n/2)^{-1}|r(\theta_n, \hat{g}_n) - \tilde{r}(\theta_n, \hat{g}_n)| \leq A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$. Before that, note that if \hat{U}_n is unbiased then the second term in Equation 18 vanishes and we have the result stated in Corollary 2.1.

Now, note that $2n^{-1} \tilde{r}(\theta_n, \hat{g}_n)$ equals

$$\log \text{IF}_{\theta_n}(\hat{\theta}_n) + \mathbb{E}_{\theta_n} \left[\log \left(1 + \frac{\hat{c}(X_n) - \text{IF}_{\theta_n}(\hat{\theta}_n)}{\text{IF}_{\theta_n}(\hat{\theta}_n)} \right) \right] - \frac{\text{Bias}_{\theta_n}(\hat{c}(X_n))}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))}$$

and using the inequality $\log(1+x) \leq x$ for all $x > -1$ on the second term on the right hand side it follows that

$$\begin{aligned}
2n^{-1} \tilde{r}(\theta_n, \hat{g}_n) &\leq \log \text{IF}_{\theta_n}(\hat{\theta}_n) + \frac{\text{Bias}_{\theta_n}(\hat{c}(X_n))}{\text{IF}_{\theta_n}(\hat{\theta}_n)} - \frac{\text{Bias}_{\theta_n}(\hat{c}(X_n))}{\mathbb{E}_{\theta_n}(\hat{c}(X_n))} \\
&= \log \text{IF}_{\theta_n}(\hat{\theta}_n) + \frac{\text{Bias}_{\theta_n}^2(\hat{c}(X_n))}{\text{IF}_{\theta_n}(\hat{\theta}_n) \mathbb{E}_{\theta_n}(\hat{c}(X_n))} = B_n.
\end{aligned}$$

Now, we write $2n^{-1}\{r(\theta_n, \hat{g}_n) - \tilde{r}(\theta_n, \hat{g}_n)\} = H_{\theta_n}(\hat{\theta}_n, \hat{c}_n) + J_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ where,

$$\begin{aligned}
H_{\theta_n}(\hat{\theta}_n, \hat{c}_n) &= \text{IF}_{\theta_n}(\hat{\theta}_n) \left\{ \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}_n} \right) - \frac{1}{\mathbb{E}_{\theta_n} \hat{c}_n} \right\} \quad \text{and} \\
J_{\theta_n}(\hat{\theta}_n, \hat{c}_n) &= (nr)^{-1} \cdot \mathbb{E}_{\theta_n} \left[\|\hat{\theta}_n - \theta_n\|^2 \left\{ \frac{1}{\hat{c}_n} - \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}_n} \right) \right\} \right].
\end{aligned}$$

Note that the second term in $H_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ can be rewritten as,

$$\mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}_n} \right) - \frac{1}{\mathbb{E}_{\theta_n} \hat{c}_n} = \frac{-1}{\mathbb{E}_{\theta_n} \hat{c}_n} \cdot \mathbb{E}_{\theta_n} \left[\left(\hat{c}_n - \mathbb{E}_{\theta_n} \hat{c}_n \right) \left(\frac{1}{\hat{c}_n} - \mathbb{E}_{\theta_n} \left(\frac{1}{\hat{c}_n} \right) \right) \right]$$

which by Cauchy-Schwartz (C-S) inequality has lower absolute value than

$$(\mathbb{E}_{\theta_n} \hat{c}_n)^{-1} \left\{ \text{Var}_{\theta_n}(\hat{c}_n) \times \text{Var}_{\theta_n}(\hat{c}_n^{-1}) \right\}^{1/2}.$$

Thus, $|H_{\theta_n}(\hat{\theta}_n, \hat{c}_n)| \leq \text{IF}_{\theta_n}(\hat{\theta}_n) \mathbb{E}_{\theta_n} \hat{c}_n^{-1} \text{SD}_{\theta_n}(\hat{c}_n) \times \text{SD}_{\theta_n}(\hat{c}_n^{-1})$.

Again, rewriting $J_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ as,

$$J_{\theta_n}(\hat{\theta}_n, \hat{c}_n) = (nr)^{-1} \mathbb{E}_{\theta_n} \left[\left\{ \|\hat{\theta}_n - \theta_n\|^2 - \mathbb{E}_{\theta_n} \|\hat{\theta}_n - \theta_n\|^2 \right\} \left\{ \hat{c}_n^{-1} - \mathbb{E}_{\theta_n}(\hat{c}_n^{-1}) \right\} \right]$$

and applying C-S inequality we get

$$|J_{\theta_n}(\hat{\theta}_n, \hat{c}_n)| \leq (nr)^{-1} \text{SD}_{\theta_n} \left(\|\hat{\theta}_n - \theta_n\|^2 \right) \cdot \text{SD}_{\theta_n}(\hat{c}_n^{-1}).$$

So $(n/2)^{-1} |r(\theta_n, \hat{g}_n) - \tilde{r}(\theta_n, \hat{g}_n)| \leq |H_{\theta_n}(\hat{\theta}_n, \hat{c}_n)| + |J_{\theta_n}(\hat{\theta}_n, \hat{c}_n)| \leq A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ and we have our desired result. \square

COROLLARY 2.1. *If \hat{U}_n is an unbiased estimate of the parameter $q(\theta_n, \hat{\theta}_n)$ and $\hat{c}_n = 1 + (nr)^{-1} \hat{U}_n > 0$ almost surely, then we have,*

$$\rho_0(\theta_n, \hat{\theta}_n) \leq \rho(\theta_n, g[\hat{\theta}_n, \hat{c}_n]) \leq \frac{1}{2} \mathbb{E}_{\theta_n} \left\{ \log \left(1 + (nr)^{-1} \hat{U}_n \right) \right\} + A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)/2.$$

The corollary follows from the above Lemma. The upper bound derived here involves expectation of a statistic along with a distortion term $A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ which will be negligible under the RASL conditions. Ignoring it for the time being we can say that an upper bound is produced when $\|\hat{\theta}_n - \theta_n\|^2$ in the lower bound of Lemma 2.6 can be replaced by a good statistic. Lemma 2.6 has an upper bound based on $\text{IF}_{\theta_n}(\hat{\theta}_n)$ and next we show that the lower bound and the upper bound are fairly close.

LEMMA 2.7. *For any point estimate $\hat{\theta}_n$ and location parameter $\theta_n \in \mathbb{R}^n$ we have,*

$$\rho_0(\theta_n, \hat{\theta}_n) \geq 2^{-1} \log \text{IF}_{\theta_n}(\hat{\theta}_n) - L_{\theta_n}(\hat{\theta}_n)/2 \text{ where,}$$

$$L_{\theta_n}(\hat{\theta}_n) = (nr)^{-1} \cdot \text{SD}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2) \cdot \text{SD}_{\theta_n} \left\{ \left(1 + (nr)^{-1} \|\hat{\theta}_n - \theta_n\|^2 \right)^{-1} \right\}$$

PROOF. From Lemma 2.5 we have

$$\begin{aligned} \log \text{IF}_{\theta_n}(\hat{\theta}_n) - 2\rho_0(\theta_n, \hat{\theta}_n) &\leq \log \text{IF}_{\theta_n}(\hat{\theta}_n) - \mathbb{E}_{\theta_n} \left\{ \log \left(1 + (nr)^{-1} \|\hat{\theta}_n - \theta_n\|^2 \right) \right\} \\ &= \mathbb{E}_{\theta_n} \left\{ \log \left(1 - \frac{\bar{l}(\theta_n, \hat{\theta}_n)}{nr + \|\hat{\theta}_n - \theta_n\|^2} \right) \right\} \end{aligned}$$

where $\bar{l}(\theta_n, \hat{\theta}_n) = \|\hat{\theta}_n - \theta_n\|^2 - q(\theta_n, \hat{\theta}_n)$ and using Jensen's inequality and $\log(1+x) \leq x$ consecutively, the difference becomes

$$\begin{aligned} &\leq -\mathbb{E}_{\theta_n} \left(\frac{\bar{l}(\theta_n, \hat{\theta}_n)}{nr + \|\hat{\theta}_n - \theta_n\|^2} \right) \\ &= -\mathbb{E}_{\theta_n} \left[\bar{l}(\theta_n, \hat{\theta}_n) \cdot \left\{ \frac{1}{nr + \|\hat{\theta}_n - \theta_n\|^2} - \mathbb{E}_{\theta_n} \left(\frac{1}{nr + \|\hat{\theta}_n - \theta_n\|^2} \right) \right\} \right] \end{aligned}$$

and by applying C-S inequality the magnitude of the said difference is

$$\leq \text{SD}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2) \times \text{SD}_{\theta_n}\{(nr + \|\hat{\theta}_n - \theta_n\|^2)^{-1}\} = L_{\theta_n}(\hat{\theta}_n).$$

This completes the proof. \square

COROLLARY 2.2. *Under the conditions of Lemma 2.6 we have*

$$[i.] \ 0 \leq \rho(\theta_n, g[\hat{\theta}_n, \hat{c}_n]) - \rho_0(\theta_n, \hat{\theta}_n) \leq 2^{-1} \{L_{\theta_n}(\hat{\theta}_n) + [A + B]_{\theta_n}(\hat{\theta}_n, \hat{c}_n)\}$$

$$[ii.] \ |\rho_0(\theta_n, \hat{\theta}_n) - 2^{-1} \log IF_{\theta_n}(\hat{\theta}_n)| \leq 2^{-1} \max \{L_{\theta_n}(\hat{\theta}_n), [A + B]_{\theta_n}(\hat{\theta}_n, \hat{c}_n)\}.$$

The corollary follows directly by combining the above lemma with Lemma 2.6. It bounds the deviation of the predictive risk from a continuous, increasing function of the MSE. The RASL conditions ensure the existence of at least one candidate for the statistic \hat{U}_n such that $c(X_n) > 0$ almost surely (follows from RASL condition P3.3) and each of the associated terms $A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$, $B_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ and $L_{\theta_n}(\hat{\theta}_n)$ is of the order of $O(n^{-1})$. Hence, Theorem 1.1 follows.

PROOF OF THEOREM 1.1. Note that under the RASL conditions we have $A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$, B_{θ_n} and $L_{\theta_n}(\hat{\theta}_n)$ to be of the order of $O(n^{-1})$. Also, note that the fact that $c > 0$ almost surely is taken care in the the RASL property P3.3. \square

2.4. Violation of the RASL conditions. Based on the lower bound of 2.5 and concentrating around the expectation by using Chebyshev's inequality we have for any a in $(0, 1)$

$$\rho_0(\theta_n, \hat{\theta}_n) \geq \frac{1}{2} \log \left(1 + \frac{(1-a)q(\theta_n, \hat{\theta}_n)}{nr} \right) \left\{ 1 - \frac{a^2 \text{Var}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2)}{q^2(\theta_n, \hat{\theta}_n)} \right\}.$$

So if **P1** of the RASL condition is violated i.e for some $\theta'_n \in \mathbb{R}^n$ we have $q(\theta'_n, \hat{\theta}_n) > O(n)$ then if **P2** holds or we have $\text{Var}_{\theta'_n}(\|\hat{\theta}_n - \theta'_n\|^2) <$

$q(\theta'_n, \hat{\theta}_n)$ then $\rho_0(\theta'_n, \hat{\theta}_n) > 1/2 \log(1 + r^{-1})$ which is the minimax risk of the best invariant density estimate and so the class of density estimates in $\mathcal{G}[1]$ centered around $\hat{\theta}$ does not have any minimax estimator. Thus, we can exclude bad point estimators in most conditions (also see Equation 17).

Among the cases where RASL conditions does not hold the only exciting case is when P2 is violated but P1 holds. In those cases the asymptotic predictive entropy risk can not be characterized in closed form. A example of a point estimator of this kind is:

$$\delta_n(i) = \begin{cases} \delta_1(X_1) & \text{if } i = 1 \\ X_i & \text{if } i = 2, \dots, n \end{cases}$$

where the univariate point estimator δ_1 is given by

$$\delta_1(x) = \begin{cases} n^{1/2} (2 \log n)^{-1/2} x & \text{if } x < (2 \log n)^{1/2} \\ x & \text{if } x \geq (2 \log n)^{1/2} \end{cases}$$

2.5. Decision Theoretic implications. The asymptotic relation between the predictive risk and the mean square risk will help us in deriving oracle inequalities on the predictive risk of g_n . The bounds will be sharp enough to discuss asymptotic optimality in the class \mathcal{G} . We would first relate the class \mathcal{G} with the other decision-theoretic classes of predictive densities. Then, we would compare the predictive risk of the respective classes in unrestricted parametric spaces.

In the above context, we consider the following 6 predictive estimates:

- $\hat{\mathbf{p}}_{\mathbf{L}}$: As an representative of the class of all Linear predictive density estimates (\mathcal{L}) we choose the predictive density $g[X_n, 1 + r]$. It is the Bayes predictive density with respect to the uniform prior, has constant risk and is inadmissible in \mathcal{L} . It is the best invariant predictive strategy and is also minimax among all procedures [Liang and Barron \(2004\)](#).
- $\hat{\mathbf{p}}_{\mathbf{E}}$: We choose the James-Stein positive part plug-in predictive density estimate $g[\hat{\theta}_n^{JS+}, r]$ as a representative of \mathcal{P} . Though the positive part James-Stein estimator is inadmissible as a point estimate, it is difficult to find estimators that have significant improvements over it. And, for all practical purposes the JS+ estimator can be considered as a ‘nearly’ admissible point estimate. In that respect we can consider

$$\hat{\mathbf{p}}_{\mathbf{E}} = \mathbf{g}[\hat{\theta}^{\mathbf{JS}+}, \mathbf{r}] \text{ where } \hat{\theta}_n^{JS+} = X_n \left(1 - \frac{(n-2)}{\|X_n\|^2} \right)_+$$

as an efficient representative from the class of Plug-in predictive densities (\mathcal{P}). The subscript stands for the class of estimative (plug-in) distributions.

- $\hat{\mathbf{p}}_{\mathbf{H}}$: We consider the Bayes predictive density estimate from the harmonic prior π_H as a representative of the class of all Bayes predictive density estimates (\mathcal{B}). It is an admissible rule. As such, it also dominates \hat{p}_L Ghosh, Mergel and Datta (2008), Komaki (2001).
- Next, we consider 3 member of \mathcal{G} which we will use to compare the risk of the predictive densities from the above 3 classes.
 - $\mathbf{g}[\hat{\theta}^{\mathbf{JS}^+}, \mathbf{1} + \mathbf{r}]$: A non-linear, fixed variance predictive density estimator around the JS+ estimator. It is uniformly better than \hat{p}_L . It is also denoted by g_M .
 - $\mathbf{g}[\hat{\theta}^{\mathbf{JS}^+}]$: The optimal member in $\mathcal{G}(\hat{\theta}^{\mathbf{JS}^+})$ which we will use to compare with $\hat{p}_{\mathbb{E}}$ and \hat{p}_L .
 - $\mathbf{g}[\hat{\theta}^{\mathbf{H}}]$: The optimal member in $\mathcal{G}(\hat{\theta}^{\mathbf{H}})$. We would like to compare its performance with \hat{p}_H . Also, $g[\hat{\theta}^{\mathbf{H}}]$ is asymptotically inadmissible among the procedures in \mathcal{G} .

In Table 1 we evaluate the predictive performance of each of these density estimates on a dataset.

Oracle inequalities and Implications. Lemma 1.1 describes the predictive risk of density estimates center around $\hat{\theta}^{\mathbf{JS}}$.

PROOF OF LEMMA 1.1. The results follows from Theorem 1.1 and by using Proposition 2.6 and Exercise 2.8 of Johnstone (2012) \square

The lemma will not be useful in very very low signal-to-noise ratio. It can be used effectively when $a_n > O(n^{-1})$. Note that, we can partition the improvement in the asymptotic prediction error over \hat{p}_L in two parts.

- We first shrink the location estimate while keeping the scale unperturbed and move to a better estimate $g[\hat{\theta}^{\mathbf{JS}^+}, \mathbf{1} + r]$. Let the improvement be denoted by d_n^1 .
- Now we optimize the scale keeping the location fixed and arrive at $g[\hat{\theta}^{\mathbf{JS}^+}]$. Let the improvement be denoted by d_n^2 .

And, based on the lemma we have,

$$d_n^1 \sim \frac{1}{2} \alpha_n \text{ and } d_n^2 \sim \frac{1}{2} \log(1 - \alpha_n)^{-1} \text{ where } \alpha_n = \{(1 + a_n)(1 + r)\}^{-1}.$$

As $\alpha_n < 1$, d_n^1 , d_n^2 as well as $d_n^2 - d_n^1$ are all positive and increasing in α_n . It means we are actually making more improvement by adapting the scale than that we got by shifting location and their difference is also decreasing in both a_n and r .

Prediction error for shrinkage estimators. By shrinkage point estimators we define estimators of the form $s(X_n)X_n$ where $s(X_n)$ is an almost everywhere differentiable function. If $\|\theta_n\|^2$ were known, then spherically symmetric shrinkage estimators of the form $s(a_n)X_n$ where $a_n = \|\theta_n\|^2/n$ and $s(a_n) \leq 1$ would be efficient. Let \mathcal{S} denotes the class of normal predictive densities based on ideal point location estimators. Such an estimate satisfies the RASL condition P2 and so Lemma 2.7 can be used to calculate an optimal lower bound on the predictive risk of the family of density estimators based on S – the class of all shrinkage point estimators conditioned on a_n .

Note that by Bias-Variance decomposition the quadratic risk of the ideal point estimator $s(a_n)X_n$ is given by

$$\mathbb{E}_{\theta_n}(\|s_n X_n - \theta_n\|^2) = s_n^2 n + \bar{s}_n^2 \|\theta_n\|^2 \text{ where } \bar{s}_n = 1 - s_n \text{ and } s_n = s(a_n).$$

Based on Lemma A.3 we have $L_{\theta_n}(\hat{\theta}_n) \leq (nr)^{-2} \text{Var}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2)$ and for an estimator in S we have,

$$\begin{aligned} \text{Var}_{\theta_n}(\|s_n X_n - \theta_n\|^2) &= \text{Var}_{\theta_n}(s_n^2 \|X_n - \theta_n\|^2 + 2 s_n \bar{s}_n \langle X_n - \theta_n, \theta_n \rangle) \\ &\leq 2[s_n^4 \text{Var}_{\theta_n}(\|X_n - \theta_n\|^2) + 4 s_n^2 \bar{s}_n^2 \text{Var}_{\theta_n}(\langle X_n - \theta_n, \theta_n \rangle)] \\ &= 2 n s_n^2 [s_n^2 + 4 \bar{s}_n^2 a_n] \end{aligned}$$

which is obviously less than $O(n)$ if $\bar{s}_n^2 a_n = O(1)$. Otherwise,

$$\begin{aligned} \rho_0(\theta_n, s(a_n)X_n) &\geq 2^{-1} \mathbb{E}_{\theta_n} \log(\|s_n X_n - \theta_n\|^2 / (nr)) \\ &\geq \mathbb{E}_{\theta_n} \log |\bar{s}_n \|\theta_n\| - s_n \chi_n| / (\sqrt{nr}) \rightarrow \infty \text{ as } n \rightarrow \infty \end{aligned}$$

and thus the optimal error in \mathcal{S} is attained at $\text{IL}(\theta_n)$ as defined in Equation (11).

Dimension independent bounds. Here, we produce a dimension independent bound on $g[\hat{\theta}_n^{JS}]$ by explicitly bounding $L_{\theta_n}(\hat{\theta}_n)$, $A_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$ and $B_{\theta_n}(\hat{\theta}_n, \hat{c}_n)$. and then substituting them in Corollary 2.2.

By construction $E(\hat{c}_n^{-1}) \leq 1$ and by using Lemma A.3 we have $\text{SD}_{\theta_n}(c_n^{-1}) \leq (nr)^{-1} \text{SD}_{\theta_n}(\hat{U})$. Now, we have,

$$(20) \quad n A_{\theta_n}(\hat{\theta}_n, \hat{c}_n) \leq r^{-2} n^{-1} \text{IF}_{\theta_n}(\hat{\theta}_n^{JS}) \text{Var}_{\theta_n}(\hat{U})$$

$$(21) \quad + r^{-3/2} n^{-1} \text{SD}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2 / n) \text{SD}_{\theta_n}(\hat{U})$$

The ideal flattening coefficient $\text{IF}_{\theta_n}(\hat{\theta}_n^{JS}) \leq (1+r^{-1})$ and the other quantities

$$(22) \quad nB_{\theta_n}(\hat{\theta}_n, \hat{c}_n) \leq r^{-2} n^{-1} \text{Bias}^2(\hat{U}_n)$$

$$(23) \quad nL_{\theta_n}(\hat{\theta}_n) \leq r^{-2} n^{-1} \text{Var}_{\theta_n}(\|\hat{\theta}_n - \theta_n\|^2)$$

For the JS estimator, for each of the terms in the R.H.S has an upper bound in terms of n and a_n . We can use the following crude upper bounds depending on n only which will provide us a uniform bound over \mathbb{R}^n :

Nature of shrinkage. The plug-in estimate $\hat{p}_{\mathbb{E}}$ performs better than \hat{p}_L and $g[\hat{\theta}^{JS+}, 1+r]$ when a_n is close to 0 but gets dominated with increasing values of a_n . And, $g[\hat{\theta}^{JS+}]$ is asymptotically better than $\hat{p}_{\mathbb{E}}$ throughout. The relationship between $g[\hat{\theta}^H]$ and \hat{p}_H can not be expressed explicitly. But following [Brown, George and Xu \(2008, Theorem 1\)](#) we can express the risk of \hat{p}_H as:

$$(24) \quad \rho(\theta_n, \hat{p}_H) = \frac{1}{2} \int_{(1+r^{-1})^{-1}}^1 v^{-1} q(\theta_n/v, \hat{\theta}^H) dv$$

where $\hat{\theta}^H$ denotes the posterior mean of the Harmonic prior. Equation 24 can be used to numerically evaluate the risk of \hat{p}_H as the risk of $\hat{\theta}_H$ has closed form. The fact that these estimators are spherically symmetric will also help. We also get the following crude bound

$$C \inf_{\beta_n \in A(\theta_n)} q(\theta_n/v, \hat{\theta}^H) \leq \rho(\theta_n, \hat{p}_H) \leq C \sup_{\beta_n \in A(\theta_n)} q(\theta_n/v, \hat{\theta}^H)$$

where $A(\theta_n) = \{ \beta_n = k \theta_n : 1 \leq k \leq \sqrt{1+r^{-1}} \}$ and $C = \log(1+r^{-1})/2$.

Minimaxity over Unrestricted Spaces. For any dimension n , \hat{p}_L is a minimax estimator. However, in dimensions greater than 2, \hat{p}_L is inadmissible and so there exists improved minimax estimators. \hat{p}_H is an improved minimax estimator than \hat{p}_L for $n \geq 3$. $g[\hat{\theta}^{JS+}]$ is also an asymptotic minimax estimator and with huge improvements over \hat{p}_L which can also be explicitly quantified. Using Theorem 1.1 asymptotically minimax predictive density estimates can be constructed around asymptotic minimax location estimates.

2.6. *A Motivational Example from Sports Betting.* Consider a game in which the outcomes depend on the actions of n players. Bets can be placed on a countable collection of (possibly overlapping) measurable sets $\mathcal{A} =$

$\{A_i : i = 1, \dots, k\}$ with $k \leq \infty$ in \mathbb{R}^n . The maximum growth rate in such a betting market is given by

$$(25) \quad \max_{Q \in \mathcal{P}(\mathbb{R}^n)} \sum_{i=1}^k P(A_i) \log \{P(A_i)/Q(A_i)\}$$

where P is the true probability distribution of the actions in the game, $\mathcal{P}(\mathbb{R}^n)$ is the set of all probability measures on \mathbb{R}^n and k is the cardinality of the collection \mathcal{A} .

Assume initially that the collection is exhaustive, i.e. $\cup_i A_i = \mathbb{R}^n$. We can construct a mutually disjoint partition $\mathcal{B} = \{B_i : 1 \leq i \leq 2^k\}$ of the collection \mathcal{A} where $B_i = \cap_{j=1}^k A_j^{w[i,j]}$ where $w[i, j]$ is the j^{th} term in the binary expansion of i and for any set $A^0 = A^c$ and $A^1 = A$. We do not track null B_i in \mathcal{B} and would ignore them through out. Let $\kappa(B_i)$ denotes the number of repetitions of the subset B_i in the collection \mathcal{A} i.e $\kappa(B_i) = \text{card}\{j : B_i \cap A_j \neq \emptyset \text{ and } j = 1, \dots, k\}$. Note that $\kappa(B_i) \in [1, k]$ and under finite overlaps we can assume that $\sup_{i=1}^{2^k} \kappa(B_i) = c < \infty$ and we define a weight function on \mathbb{R}^n as $w(x) = \sum_{i=1}^{2^k} c^{-1} \kappa(B_i) \mathbb{I}_{B_i}(x)$. Note that, $w(x) \in (0, 1]$ acts as a tilt function for the densities $p(x)$ and $q(x)$.

THEOREM 2.1. *If the probability measure P and Q have densities p and q with respect to Lebesgue measure, then for any countable collection of exhaustive measurable sets \mathcal{A} we have,*

$$\sum_{i=1}^k P(A_i) \log \{P(A_i)/Q(A_i)\} \leq c \cdot D(p||q)$$

where $D(p||q) = \int p(x) \log \{p(x)/q(x)\} dx$ is the differential relative entropy between P and Q .

PROOF. If the collection consists of mutually disjoint sets then the proof follows from the data processing inequalities associated with quantization idea in information theory. The function $t \log t$ is strictly convex if $t > 0$. So for any positive random variable T and any sigma-finite measure, by Jensen's inequality we have, $E_\mu(T \log T) \geq E_\mu(T) \log E_\mu(T)$. For any measurable set A , with $T(x) = p(x)/q(x)$ and measure $\mu(x) = q(x)/Q(A) dx$ we have $P(A) \log P(A)/Q(A) \leq \int_A p(x) \log \{p(x)/q(x)\} dx$ and so the proof extends to mutually exclusive cases.

If the events are not mutually disjoint then we can construct its mutually disjoint partition $\mathcal{B} = \{B_i : 1 \leq i \leq 2^k\}$ as above and using the Log-Sum

inequality (Cover and Thomas, 1991, Theorem 2.7.1) separately on each A_i we have,

$$\sum_{i=1}^k P(A_i) \log\{P(A_i)/Q(A_i)\} \leq \sum_{i=1}^{2^k} \kappa(B_i) P(B_i) \log\{P(B_i)/Q(B_i)\}$$

and again using the above quantization argument we can show that the R.H.S above is less than $c \int w(x)p(x) \log\{p(x)/q(x)\} dx \triangleq c D(w.p||w.q)$. Now, observe that

$$\begin{aligned} D(w.p||w.q) - D(p||q) &= \int (1 - w(x)) p(x) \log\{q(x)/p(x)\} dx \\ &\leq \log \left[\int (1 - w(x)) q(x) dx \right] \end{aligned}$$

by Jensen's inequality and the result follows as $\int (1 - w(x)) q(x) dx \leq 1$. \square

If the collection is not exhaustive we can restrict our densities to the corresponding subsets of \mathbb{R}^n .

2.6.1. An illustration with a Dataset. We consider the Baseball data that was used to show the advantage of shrinking location estimates in Efron and Morris (1977). The dataset consists of 18 players (so, $n = 18$ which is not so high dimensions) with exactly 45 at-bats on a particular date during the 1970 season. The objective is to predict the performance of the players on the remainder of the season .

The number of hits (H) and the number of at-bats (N) over two portions of the season were

$$H_{ji} \stackrel{ind.}{\sim} \text{Binomial}(N_{ji}, p_i), \quad j = 1, 2; \quad i = 1, \dots, n.$$

Where $j = 1$ denotes past data and $j = 2$ represents the unknown future. As the variance of the Binomial model depends of the mean parameter p_i , a variance stabilization transformation Brown (2008) is conducted (which goes through as N_{ij} are quite large). The transformation

$$(26) \quad X_{ji} = \arcsin \left(\frac{H_{ji} + 1/4}{N_{ji} + 1/2} \right)^{1/2}$$

reduces the binomial model to the normal model

$$(27) \quad X_{ji} \sim N(\theta_i, \sigma_{ji}^2) \quad \text{where } \theta_i = \arcsin \sqrt{p_i}, \quad \sigma_{ij}^2 = (4N_{ji})^{-1}$$

r	\hat{p}_E	\hat{p}_L	g_M	$g[\hat{\theta}^{JS+}]$	$g[\hat{\theta}^H]$	\hat{p}_H
0.1	22.963	19.451	15.487	11.435	19.232	19.578
0.2	11.482	14.174	10.539	7.418	13.982	14.289
0.5	4.593	8.326	5.418	3.717	8.188	8.424
1	2.296	5.067	2.886	2.047	4.975	5.142
2	1.148	2.868	1.415	1.081	2.815	2.924
5	0.459	1.250	0.524	0.448	1.227	1.286
10	0.23	0.645	0.248	0.227	0.633	0.614

TABLE 1

Predictive loss of different Gaussian strategies on the Baseball data.

and X_i independent for $1 \leq i \leq n$. With the past $P = X_1$ and the future $F = X_2$, we have the following predictive set-up :

$$(28) \quad F|\theta_n \sim N(\theta_n, v_y I_n) ; \quad P|\theta_n \sim N(\theta_n, v_x I_n).$$

We want joint predictive densities of the future performances of players in this standardized model. We use a very naive evaluation strategy by considering the entire season's batting average as the true parametric value. In the entire season the players ended up playing around 400 games on the average. So, evaluating the predictive densities at $\theta_i^0 = \arcsin(p_i^{\text{full}})^{1/2}$ where p_i^{full} are the batting averages from the entire season will not be terrible. Evaluation procedures with guarantees may be developed in a sequential set-up [Lai, Gross and Shen \(2011\)](#).

While using shrinkage on the location estimators we shrink towards the grand average. We evaluate the 6 different predictive strategies of Section 2.5 for different values of the future to past variability. The value of r will be close to 0.1 when we consider prediction on the entire remaining half of the season.

We find that for any choice r , $g[\hat{\theta}^{JS+}]$ is the best one among the 6 estimators considered. Also, $d_n^2 - d_n^1$ (as discussed in Section 2.5) is decreasing in r . \hat{p}_E behaves well when r is large and horribly for small values. The losses for \hat{p}_H and $g[\hat{\theta}^H]$ are very similar.

3. Restricted minimax predictive risk of $\mathcal{G}[p]$.

A typical member in the class $\mathcal{G}[p]$ of all product Gaussian predictive densities is represented by $g[\hat{\theta}_n, \hat{D}_n] = \prod_{i=1}^n N(\hat{\theta}(i), \hat{d}(i) \sigma_f^2)$. Generalizing the

argument in Lemma 2.5 we see that a lower bound on the minimum predictive risk $\rho_p(\theta_n, \hat{\theta}_n)$ of all density estimates in $\mathcal{G}_n[p]$ that have mean $\hat{\theta}_n$, is given by

$$(29) \quad \rho_p(\theta_n, \hat{\theta}_n) \geq \frac{1}{2} \sum_{i=1}^n \mathbb{E}_{\theta(i)} \{ \log(1 + r^{-1} (\hat{\theta}(i) - \theta(i))^2) \}.$$

The predictive risk of the estimate $g[\hat{\theta}_n, \hat{D}_n]$ is given by

$$(30) \quad 2\rho(\theta_n, g[\hat{\theta}_n, \hat{D}_n]) = \sum_{i=1}^n \mathbb{E}_{\theta_n} \log(\hat{d}(i)) + \mathbb{E}_{\theta_n} \sum_{i=1}^n \left\{ \frac{1 + (\hat{\theta}(i) - \theta(i))^2}{\hat{d}(i)} - \hat{d}(i) \right\}.$$

It is not necessarily true that

$$\rho_p(\theta_n, \hat{\theta}_n) = \min_{\hat{D}_n \in \mathbb{R}_+^n} \rho(\theta_n, g[\hat{\theta}_n, \hat{D}_n])$$

asymptotically equals the lower bound given in Equation (29). In the previous section we saw that under sufficient regularity conditions these bounds matches. The ideas there can be extended to block-wise estimators and to non-orthogonal models by using the concept of Mallows's unbiased risk estimates. In the ℓ_0 sparse predictive space as the degree of sparsity tends to zero, i.e., $s/n \rightarrow 0$ as $n \rightarrow \infty$, the lower bound given in Equation (29) is significantly greater than the minimax predictive risk over $\mathcal{G}[p]$. And so, procedure used in the previous section can not be used for finding the asymptotic minimax predictive Gaussian risk over $\Theta(n, s)$.

Minimax predictive risk over sparse parameter spaces. Here we outline the proof of Theorem 1.3. Following the Bayes-Minimax procedure of Johnstone (2012) (Chapter 4.4) the multivariate minimax problem can be reduced to univariate minimax problem with moment prior constraints

$$\mathfrak{m}(\eta) = \{\pi \in \mathcal{P}(\mathbb{R}) : \pi(0) \geq 1 - \eta\}$$

where $\mathcal{P}(\mathbb{R})$ is the collection of all probability measures on \mathbb{R} . In Theorem 1.1 in Mukherjee and Johnstone (2012) we have the univariate minimax risk

$$\min_{\hat{p}} \max_{\pi \in \mathfrak{m}(\eta)} \int \rho(\theta, \hat{p}) \pi(\theta) d\theta \sim (1 + r)^{-1} \eta \log \eta^{-1} \text{ as } \eta \rightarrow 0.$$

When restricted to the Gaussian family the minimax risk will be

$$\min_{\hat{p} \in \mathcal{G}} \max_{\pi \in \mathfrak{m}(\eta)} \int \rho(\theta, \hat{p}) \pi(\theta) d\theta \sim f(\eta) \text{ as } \eta \rightarrow 0$$

where $f(\eta) = r^{-1} \eta \log \eta^{-1}$. In this univariate asymptotic set-up the lower bound in Equation (29) is much lower than the asymptotic rate $\eta \log \eta^{-1}$ and hence unusable. We get an upper bound on the minimax Gaussian risk as from point estimation theory [Donoho and Johnstone \(1994\)](#) it follows that the minimax plug-in risk in this asymptotic set-up is $f(\eta)$. For a lower bound consider the predictive risk of the normal density estimate $g[\hat{\theta}, \hat{d}]$

$$(31) \quad \rho(\theta, g[\hat{\theta}, \hat{d}]) = \mathbb{E}_\theta(\log \hat{d}) + \mathbb{E}_\theta\{\hat{d}^{-1} \cdot (1 + (\hat{\theta} - \theta)^2) - 1\}.$$

And the idea is to establish the necessity of threshold zone as done in [Johnstone and Silverman \(2004\)](#). For $\rho(0, g[\hat{\theta}, \hat{d}])$ – the predictive risk of $g[\hat{\theta}, \hat{d}]$ at the origin, to be lower than the order of η we need a threshold size of at least $\lambda(\eta) = \sqrt{2 \log \eta^{-1}}$. And for density estimators of the form

$$(32) \quad \hat{p}[\lambda(\eta)](\cdot|X) = \begin{cases} N(0, \sigma_f^2) & \text{if } |X| \leq \lambda(\eta) \\ N(\hat{\theta}(X), \hat{d}(X) \sigma_f^2) & \text{if } |X| > \lambda(\eta) \end{cases}$$

the supremum predictive risk at the non-zero support points is $f(\eta)$, i.e.,

$$(33) \quad \sup_{\theta \neq 0} \rho(\theta, \hat{p}[\lambda(\eta)]) \sim f(\eta) \text{ as } \eta \rightarrow 0.$$

Thus, it follows that sup-optimality of the class $\mathcal{G}[p]$ is $1 + r^{-1}$.

Acknowledgements.

G.M. is indebted to Thomas Cover for the numerous enjoyable conversations, and his insights on information complexity and density estimation in predictive context, which inspired and reshaped many ideas in the paper.

APPENDIX A

LEMMA A.1. *Y_n is sequence of random variables such that $Y_n \stackrel{d}{=} \chi_n^2(\lambda_n)$ for a non-negative and increasing sequence $\{\lambda_n : n \geq 1\}$ then for $n \geq 5$ we have*

$$\text{Var}(Y_n^{-1}) \leq k_1(n) \cdot n^{-3} \text{ where } k_1(n) = 3(1 - 2/n)^{-2}(1 - 4/n)^{-1}.$$

PROOF. We observe that Y_n being a non-central chi-square random variable can be written as convolution of central Chi-square and Poisson random variables

$$Y_n \stackrel{d}{=} \chi_{n+2N}^2 \text{ where } N_n \stackrel{d}{=} \text{Poisson}(\lambda_n/2).$$

Decomposing the variance by conditioning on the Poisson random variable we have,

$$\begin{aligned}\mathbb{V}ar(Y_n^{-1}) &= \mathbb{V}ar_{\lambda_n}\left(E(Y_n^{-1}|N_n)\right) + E_{\lambda_n}\left(\mathbb{V}ar(Y_n^{-1}|N_n)\right) \\ &= \mathbb{V}ar_{\lambda_n}\left(\frac{1}{n+2N_n-2}\right) + E_{\lambda_n}\left(\frac{2}{(n+2N_n-2)^2(n+2N_n-4)}\right)\end{aligned}$$

which follows from moments of central chi-square (gamma) distribution and as $N_n \geq 0$ the second term on the R.H.S is $\leq 2(n-2)^{-2}(n-4)^{-1}$ and by Lemma A.3 we have

$$\begin{aligned}(n-2)^2 \mathbb{V}ar_{\lambda_n}\left(\frac{1}{n+2N_n-2}\right) &= \mathbb{V}ar_{\lambda_n}\left(\frac{1}{1+2N_n/(n-2)}\right) \\ &\leq \{1+2E(N_n)/(n-2)\}^{-4} \mathbb{V}ar\left(\frac{2N_n}{n-2}\right) \\ &= \frac{4\lambda_n(n-2)^2}{(n-2+2\lambda_n)^4} \leq \frac{1}{2(n-2)}.\end{aligned}$$

Thus, $\mathbb{V}ar(Y_n^{-1}) \leq 3(n-2)^{-2}(n-4)^{-1}$. \square

LEMMA A.2. *If $Y_n \stackrel{d}{=} \chi_n^2(\lambda_n)$ and λ_n is an increasing sequence then*

$$\lambda_n^2 P(Y_n \leq n-2) \leq O(n)$$

PROOF. Holds trivially for $\lambda_n \leq O(\sqrt{n})$. So we will prove for all other sequences i.e sequence where λ_n/\sqrt{n} is not bounded. Note that $P(Y_n \leq n-2) \leq P(Y_n \leq n)$. And as Y_n is a non-central chi-square we have

$$Y_n \stackrel{d}{=} V_{n+2N} \text{ where } N \stackrel{d}{=} \text{Poisson}(\lambda_n) \text{ and } V_n \stackrel{d}{=} \chi_n^2(0)$$

Now, for any fixed n and N we have,

$$P(V_{n+2N} \leq n) \leq 2P(V_{m+2N} \leq m) \text{ for all } m \geq n \text{ such that } m-n \text{ is large.}$$

Because $P(V_{m+2N} \leq m|V_{n+2N} \leq n) \leq P(\chi_{m-n}^2(0) \leq m-n) \leq 1/2$. So,

$$\begin{aligned}\lim_{n \rightarrow \infty} P(Y_n \leq n) &= \lim_{n \rightarrow \infty} E_{\lambda_n} \left\{ P(V_{n+2N} \leq n | N) \right\} \\ &\leq 2 \lim_{n \rightarrow \infty} E_{\lambda_n} \left\{ \lim_{n \rightarrow \infty} P(V_{n+2N} \leq n | N) \right\} \\ &= 2 \lim_{n \rightarrow \infty} E_{\lambda_n} \left\{ \Phi\left(\frac{-2N}{\sqrt{2n+4N}}\right) \right\}\end{aligned}$$

as we can do a normal approximation to the sequence of central chi-square random variables. Next, we interchange the integrals (by Fubini's as integrand is positive) and then use bounded convergence theorem to have,

$$\begin{aligned} \lim_{n \rightarrow \infty} P(Y_n \leq n) &\leq 2 \lim_{n \rightarrow \infty} \int \phi(z) P_{\lambda_n} \left(\frac{2N}{\sqrt{2n+4N}} \leq z \right) dz \\ &= 2 \int \phi(z) \lim_{n \rightarrow \infty} P_{\lambda_n} \left(\frac{2N}{\sqrt{2n+4N}} \leq z \right) dz \end{aligned}$$

Now for all large n , λ_n is large (as λ_n increasing and λ_n/\sqrt{n} is not a bounded sequence). So each large n , we can separately do a normal approximation to the Poisson random variable N .

Consider the case first when $\lambda_n > O(n)$. In this case the following naive bound will work:

$$P_{\lambda_n} \left(\frac{2N}{\sqrt{2n+4N}} \leq z \right) \leq P_{\lambda_n} \left(\frac{\sqrt{N}}{\sqrt{n}} \leq z \right) \sim \tilde{\Phi} \left(\frac{\lambda_n - nz^2}{\sqrt{\lambda_n}} \right).$$

We will use this bound for all z such that $z^2 \leq t_n$ where t_n equals $n^{-1}(\lambda_n - \sqrt{\lambda_n}\sqrt{4\log \lambda_n + 2\log n})$. Also note that,

$$\lambda_n^2 \tilde{\Phi} \left(\frac{\lambda_n - nz^2}{\sqrt{\lambda_n}} \right) \leq O(n) \text{ for all } z^2 \leq t_n \text{ and } \tilde{\Phi}(t_n) = O(n\lambda_n^{-2}).$$

And so, it follows that $\lambda_n^2 \lim_{n \rightarrow \infty} P(Y_n \leq n) \leq O(n)$. \square

LEMMA A.3. *For any non-negative random variable Y*

$$\mathbb{V}ar\{(1+Y)^{-1}\} \leq \{1+E(Y)\}^{-4} \mathbb{V}ar(Y).$$

PROOF. As Y is non-negative we have

$$\left(\frac{1}{1+Y} - \frac{1}{1+E(Y)} \right)^2 = \frac{(Y - E(Y))^2}{(1+Y)^2(1+EY)^2} \leq \frac{(Y - E(Y))^2}{(1+EY)^2}.$$

Now, taking expectation on both sides and using Bias-Variance decomposition we get

$$\mathbb{V}ar \left(\frac{1}{1+Y} \right) + \left(E \left(\frac{1}{1+Y} \right) - \frac{1}{1+E(Y)} \right)^2 \leq \{1+E(Y)\}^{-4} \mathbb{V}ar(Y).$$

This completes the proof. \square

LEMMA A.4. *For any random variable X we have $\mathbb{V}ar(X_+) \leq \mathbb{V}ar(X)$.*

PROOF. With the decomposition of $X = X_+ - X_-$ we have

$$\begin{aligned}\text{Var}(X) &= \text{E}(X^2) - \text{E}^2(X) \\ &= \text{E}(X_+^2) + \text{E}(X_-^2) - \text{E}^2(X_+) - \text{E}^2(X_-) + 2\text{E}(X_+)\text{E}(X_-) \\ &= \text{Var}(X_+) + \text{Var}(X_-) + 2\text{E}(X_+)\text{E}(X_-)\end{aligned}$$

and we get the stated result as all the terms in R.H.S. are non-negative. \square

REFERENCES

- AITCHISON, J. (1975). Goodness of prediction fit. *Biometrika* **62** 547–554. [MR0391353 \(52 #12174\)](#)
- ASLAN, M. (2006). Asymptotically minimax Bayes predictive densities. *Ann. Statist.* **34** 2921–2938. [MR2329473 \(2008g:62093\)](#)
- BARRON, A., RISSANEN, J. and YU, B. (1998). The minimum description length principle in coding and modeling. *IEEE Trans. Inform. Theory* **44** 2743–2760. Information theory: 1948–1998. [MR1658898 \(99h:94032\)](#)
- BELL, R. M. and COVER, T. M. (1980). Competitive optimality of logarithmic investment. *Math. Oper. Res.* **5** 161–166. [MR571810 \(81g:90114\)](#)
- BJØRNSTAD, J. F. (1990). Predictive likelihood: a review. *Statist. Sci.* **5** 242–265. With comments and a rejoinder by the author. [MR1062578 \(91f:62002\)](#)
- BRODIEA, J., DAUBECHIES, I., MOLC, C. D., GIANNONED, D. and LORISC, I. (2009). Sparse and stable Markowitz portfolios. *Proceedings of National Academy of Sciences*.
- BROWN, L. D. (1971). Admissible estimators, recurrent diffusions, and insoluble boundary value problems. *Ann. Math. Statist.* **42** 855–903. [MR0286209 \(44 #3423\)](#)
- BROWN, L. D. (2008). In-season prediction of batting averages: a field test of empirical Bayes and Bayes methodologies. *Ann. Appl. Stat.* **2** 113–152. [MR2415597](#)
- BROWN, L. D., GEORGE, E. I. and XU, X. (2008). Admissible predictive density estimation. *Ann. Statist.* **36** 1156–1170. [MR2418653 \(2009i:62023\)](#)
- BUTLER, R. W. (1986). Predictive likelihood inference with applications. *J. Roy. Statist. Soc. Ser. B* **48** 1–38. With discussion and a reply by the author. [MR848048 \(87k:62001\)](#)
- COVER, T. M. and THOMAS, J. A. (1991). *Elements of information theory*. Wiley-Interscience, New York, NY, USA.
- DEMBO, A. and ZEITOUNI, O. (1993). *Large deviations techniques and applications*. Jones and Bartlett Publishers, Boston, MA. [MR1202429 \(95a:60034\)](#)
- DONOHU, D. L. (2006). Compressed sensing. *IEEE Trans. Inform. Theory* **52** 1289–1306. [MR2241189 \(2007e:94013\)](#)
- DONOHU, D. L. and JOHNSTONE, I. M. (1994). Minimax risk over l_p -balls for l_q -error. *Probab. Theory Related Fields* **99** 277–303. [MR1278886 \(95g:62019\)](#)
- DONOHU, D., JOHNSTONE, I. and MONTANARI, A. (2011). Accurate Prediction of Phase Transitions in Compressed Sensing via a Connection to Minimax Denoising. [arXiv:1111.1041](#).
- EFRON, B. (2004). The estimation of prediction error: covariance penalties and cross-validation. *J. Amer. Statist. Assoc.* **99** 619–642. With comments and a rejoinder by the author. [MR2090899](#)
- EFRON, B. (2010). *Large-scale inference. Institute of Mathematical Statistics (IMS) Monographs* **1**. Cambridge University Press, Cambridge. Empirical Bayes methods for estimation, testing, and prediction. [MR2724758 \(2012a:62006\)](#)
- EFRON, B. and MORRIS, C. (1977). Stein’s Paradox in Statistics. *Scientific American*.

- FOURDRINIER, D., MARCHAND, É., RIGHI, A. and STRAWDERMAN, W. E. (2011). On improved predictive density estimation with parametric constraints. *Electron. J. Stat.* **5** 172–191. . [MR2792550 \(2012h:62088\)](#)
- FRIEDMAN, J., HASTIE, T. and TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- GEISSER, S. (1971). The inferential use of predictive distributions. In *Foundations of statistical inference (Proc. Sympos., Univ. Waterloo, Ont., 1970)* 456–469. Holt, Rinehart and Winston of Canada, Toronto, Ont. With comments by V. P. Godambe, I. J. Good, W. J. Hall and D. A. Sprott and a reply by the author. [MR0381054 \(52 ##1951\)](#)
- GEORGE, E. I., LIANG, F. and XU, X. (2006). Improved minimax predictive densities under Kullback-Leibler loss. *Ann. Statist.* **34** 78–91. . [MR2275235 \(2008h:62034\)](#)
- GEORGE, E. I., LIANG, F. and XU, X. (2012). From minimax shrinkage estimation to minimax shrinkage prediction. *Statist. Sci.* **27** 82–94. . [MR2953497](#)
- GEORGE, E. I. and XU, X. (2008). Predictive density estimation for multiple regression. *Econometric Theory* **24** 528–544. . [MR2391619 \(2009a:62121\)](#)
- GHOSH, M., MERGEL, V. and DATTA, G. S. (2008). Estimation, prediction and the Stein phenomenon under divergence loss. *J. Multivariate Anal.* **99** 1941–1961. . [MR2466545 \(2009m:62027\)](#)
- HARTIGAN, J. A. (1998). The maximum likelihood prior. *Ann. Statist.* **26** 2083–2103. . [MR1700222 \(2000h:62030\)](#)
- HINKLEY, D. (1979). Predictive likelihood. *Ann. Statist.* **7** 718–728. [MR532237 \(81m:62007a\)](#)
- JOHNSTONE, I. M. (2012). Gaussian Estimation: Sequence and Wavelet Models. Available at: "<http://www-stat.stanford.edu/~imj>".
- JOHNSTONE, I. M. and SILVERMAN, B. W. (2004). Needles and straw in haystacks: empirical Bayes estimates of possibly sparse sequences. *Ann. Statist.* **32** 1594–1649. . [MR2089135 \(2005h:62027\)](#)
- KOMAKI, F. (1996). On asymptotic properties of predictive distributions. *Biometrika* **83** 299–313. . [MR1439785 \(98d:62048\)](#)
- KOMAKI, F. (2001). A shrinkage predictive distribution for multivariate normal observables. *Biometrika* **88** 859–864. . [MR1859415](#)
- KUBOKAWA, T., MARCHAND, E., STRAWDERMAN, W. E. and TURCOTTE, J. P. (2012). Minimality in Predictive Density Estimation with Parametric Constraints. *Technical Report, available at* <http://www.usherbrooke.ca/mathematiques/recherche/publications/rapports-recherche/>.
- KULLBACK, S. and LEIBLER, R. A. (1951). On information and sufficiency. *Ann. Math. Statistics* **22** 79–86. [MR0039968 \(12,623a\)](#)
- LAI, T. L., GROSS, S. T. and SHEN, D. B. (2011). Evaluating probability forecasts. *Ann. Statist.* **39** 2356–2382. . [MR2906871](#)
- LARIMORE, W. E. (1983). Predictive inference, sufficiency, entropy and an asymptotic likelihood principle. *Biometrika* **70** 175–181. . [MR742987 \(85k:62005\)](#)
- LAURITZEN, S. L. (1974). Sufficiency, prediction and extreme models. *Scand. J. Statist.* **1** 128–134. [MR0378162 \(51 ##14331\)](#)
- LEDoux, M. (2001). *The concentration of measure phenomenon. Mathematical Surveys and Monographs* **89**. American Mathematical Society, Providence, RI. [MR1849347 \(2003k:28019\)](#)
- LIANG, F. and BARRON, A. (2004). Exact minimax strategies for predictive density estimation, data compression, and model selection. *IEEE Trans. Inform. Theory* **50** 2708–2726. . [MR2096988 \(2005f:94040\)](#)

- LIANG, F. and BARRON, A. (2005). *Exact Minimax Predictive Density Estimation and MDL*. *Advances in Minimum Description Length: Theory and Applications* (P. Grunwald, I. Myung and M. Pitt eds). MIT Press.
- McMILLAN, B. (1956). Two inequalities implied by unique decipherability. *Information Theory, IRE Transactions on* **2** 115–116.
- MUKHERJEE, G. and JOHNSTONE, I. M. (2012). Exact minimax estimation of the predictive density in sparse Gaussian models. arXiv:1211.2071.
- NUSSBAUM, M. (1996). Asymptotic equivalence of density estimation and Gaussian white noise. *Ann. Statist.* **24** 2399–2430. . [MR1425959 \(98k:62065\)](#)
- RISSANEN, J. (1984). Universal coding, information, prediction, and estimation. *IEEE Trans. Inform. Theory* **30** 629–636. . [MR755791 \(85g:94009\)](#)
- ROBBINS, H. (1956). An empirical Bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, 1954–1955, vol. I* 157–163. University of California Press, Berkeley and Los Angeles. [MR0084919 \(18,947e\)](#)
- STEIN, C. (1974). Estimation of the mean of a multivariate normal distribution. In *Proceedings of the Prague Symposium on Asymptotic Statistics* (Charles Univ., Prague, 1973), Vol. II 345–381. Charles Univ., Prague. [MR0381062 \(52 #1959\)](#)
- STEIN, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *Ann. Statist.* **9** 1135–1151. [MR630098 \(83a:62080\)](#)
- TIBSHIRANI, R. (2011). Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*.
- TIBSHIRANI, R., HASTIE, T., NARASIMHAN, B. and CHU, G. (2002). Diagnosis of multiple cancer types by shrunken centroids of gene expression. *Proceedings of National Academy of Sciences*.
- XU, J. (2007). A closed form for the harmonic-prior Bayes estimator with associated confidence sets for the mean of a multivariate normal distribution PhD thesis "<http://repository.upenn.edu/dissertations/AAI3271836>".
- XU, X. and LIANG, F. (2010). Asymptotic minimax risk of predictive density estimation for non-parametric regression. *Bernoulli* **16** 543–560. . [MR2668914 \(2011g:62110\)](#)
- XU, X. and ZHOU, D. (2011). Empirical Bayes predictive densities for high-dimensional normal models. *J. Multivariate Analysis* **102** 1417–1428.
- YANG, Y. (2007). Consistency of cross validation for comparing regression procedures. *Ann. Statist.* **35** 2450–2473. . [MR2382654 \(2009a:62022\)](#)
- YE, J. (1998). On measuring and correcting the effects of data mining and model selection. *J. Amer. Statist. Assoc.* **93** 120–131. . [MR1614596 \(99a:62097\)](#)

ADDRESS OF THE FIRST AND SECOND AUTHORS
DEPARTMENT OF STATISTICS
SEQUOIA HALL, 390 SERRA MALL
STANFORD UNIVERSITY
STANFORD, CA 94305-4065
E-MAIL: gourab@stanford.edu
imj@stanford.edu